

Time optimal trajectories for bounded velocity differential drive vehicles

Devin J. Balkcom Matthew T. Mason

November 16, 2006

Abstract

This paper presents the time optimal trajectories for differential drive vehicles in the unobstructed plane. The wheel angular velocities are bounded, but may be discontinuous. The paper proves the existence of optimal controls, derives the structure of optimal trajectories, and develops an algorithm for producing a time optimal trajectory between any two configurations. Every nontrivial optimal trajectory is composed of straight segments alternating with turns about the robot's center. Optimal trajectories may have as many as five actions, but four actions are sufficient—for every optimal trajectory of five actions, there is an equally fast trajectory with four actions.

1 Introduction

This paper derives the time optimal trajectories for a bounded velocity differential drive robot in the unobstructed plane. “Differential drive” means the robot is like a conventional wheelchair, using two independently driven wheels to maneuver in the plane. “Bounded velocity” means that the wheel angular velocities are bounded, but may be discontinuous. Under these assumptions, we will see that time optimal trajectories exist for all choices of start and goal, and are composed of straight lines alternating with turns about the center of the robot. Optimal trajectories contain at most three straight segments and two turns. There are a number of other restrictions, leading to a set of 40 different combinations arranged in 9 different symmetry classes. The simplest motion to a generic




Figure 1: A time optimal trajectory from each non-trivial symmetry class. Classes H and I are only valid for small motions—they are magnified five times and the robot outline is not drawn.

goal is a turn-straight-turn: turn to face the goal, roll straight forwards to the goal, and turn to the goal orientation. In more interesting cases, the optimal trajectory contains a turn at an intermediate “via” point. Figure 1 shows an example trajectory for each of the nine symmetry classes.

The key results of the paper are:

- Proof of the existence of time optimal trajectories, and an analysis of the structure of the time optimal trajectories.
- An algorithm for determining all optimal trajectories between a start and goal pair, and for determining the time of the optimal trajectories.
- Plots of the “balls” in configuration space: sets of configurations reachable in a given time.

We apply the methodology pioneered in the work of Sussmann and Tang, Souères and Boissonat, and Laumond, to find the time optimal trajectories:

1. Apply Pontryagin’s maximum principle to obtain extremal trajectories. The optimal trajectories are a subset of the extremal trajectories.
2. Develop the properties of the switching functions to show that the extremal trajectories have a simple structure.
3. Use the geometry of the extremals to derive further necessary conditions for optimality.
4. Determine the start and goal configurations for which each candidate extremal may be optimal.

1.1 Previous work.

Most of the work on time optimal control with bounded velocity models has focused on steered vehicles rather than differential drives, originating with papers by Dubins [?] and by Reeds and Shepp [?]. Many of the techniques employed here are an extension of optimal control techniques developed for steered vehicles by Souères and Boissonat [?], Souères and Laumond [?] and Sussmann and Tang [?].




Figure 2: Notation

Previous work on differential drive robots has assumed bounds on acceleration; for example, see papers by Reister and Pin [?] for a numerical approach, and Renaud and Fourquet [?] for work on the structure of optimal trajectories. The bounded velocity model is simpler than the bounded acceleration model, and we determine the structures and costs of the fastest trajectories analytically.

2 Model, assumptions, notation

We will assume each wheel angular velocity is in the interval $[-1, 1]$. Accordingly, define a control region:

$$U = [-1, 1] \times [-1, 1] \quad (1)$$

and consider the class of *admissible controls* to be the measurable functions $u(t)$ mapping the time interval $[0, T]$ to U :

$$u(t) = \begin{pmatrix} \omega_1(t) \\ \omega_2(t) \end{pmatrix} \quad (2)$$

$$u(t) : [0, T] \rightarrow U \quad (3)$$

Let v and ω be defined

$$v = \frac{\omega_1 + \omega_2}{2} \quad (4)$$

$$\omega = \frac{\omega_2 - \omega_1}{2b} \quad (5)$$

and let the state trajectory $q(t) = (x(t), y(t), \theta(t))$ be defined for any initial state q_0 and admissible control

$u(t)$ using Lebesgue integration:

$$x(t) = x_0 + \int_0^t v \cos \theta \quad (6)$$

$$y(t) = y_0 + \int_0^t v \sin \theta \quad (7)$$

$$\theta(t) = \theta_0 + \int_0^t \omega \quad (8)$$

Throughout the paper we use Lebesgue integration assuming the standard measure.

It follows that for any admissible control $u(t)$ and initial state q_0 , an absolutely continuous state trajectory $q(t)$ exists and has a time derivative almost everywhere, with:

$$\dot{q}(t) = \begin{pmatrix} v \cos \theta \\ v \sin \theta \\ \omega \end{pmatrix} \quad \text{a.e.} \quad (9)$$

The abbreviation ‘‘a.e.’’ stands for ‘‘almost everywhere,’’ meaning that the equation holds everywhere except on some set of measure zero. We cannot say the equality holds for all t since the controls are only required to be measurable, not continuous or piecewise continuous. But two controls that are equal almost everywhere will yield equivalent trajectories. Readers unfamiliar with measure theory will not go too far wrong by assuming the controls are piecewise continuous, and ignoring the measure theory jargon. Indeed one result of this work is to show that the time optimal trajectories may be obtained using piecewise continuous controls.

The control system equations can be written in a more convenient form if we define vector fields f_1 and f_2 :

$$f_1 = \begin{pmatrix} \frac{1}{2} \cos(\theta) \\ \frac{1}{2} \sin(\theta) \\ -\frac{1}{2b} \end{pmatrix}, f_2 = \begin{pmatrix} \frac{1}{2} \cos(\theta) \\ \frac{1}{2} \sin(\theta) \\ \frac{1}{2b} \end{pmatrix} \quad (10)$$

Now the control system equations can be written:

$$\dot{q} = \omega_1 f_1 + \omega_2 f_2 \quad \text{a.e.} \quad (11)$$

The vector fields f_1 and f_2 describe two freedoms of the robot. We introduce an additional vector field,

the Lie bracket of f_1 and f_2 :

$$f_3 = [f_1, f_2] = \begin{pmatrix} \frac{1}{2b} \sin(\theta) \\ -\frac{1}{2b} \cos(\theta) \\ 0 \end{pmatrix} \quad (12)$$

The vector field f_3 describes a sideways motion that cannot be achieved by any linear combination of f_1 and f_2 , but can be approximated by parallel parking motions of the robot.

We also define rectified path length in \mathbf{E}^2 , the plane of robot positions:

$$s(t) = \int_0^t |v| \quad (13)$$

and rectified arc length in \mathbf{S}^1 , the circle of robot orientations:

$$\sigma(t) = \int_0^t |\omega| \quad (14)$$

We introduce some additional notation for trajectories. Later sections show that extremal trajectories are composed of straight lines and turns about the robot’s center. We will represent *forwards* by \uparrow , *backwards* by \downarrow , *left turn* by \curvearrowright , and *right turn* by \curvearrowleft . Thus the trajectory structure $\curvearrowleft\uparrow\curvearrowright$ can be read ‘‘left forwards left’’. The symbols \uparrow , \downarrow , \curvearrowright , and \curvearrowleft correspond to vertices of U :

ω_1	ω_2	Notation
1	1	\uparrow
-1	-1	\downarrow
-1	1	\curvearrowright
1	-1	\curvearrowleft

When necessary, a subscript will indicate the distance or angle traveled. We will use \mathbf{t} and \mathbf{s} to represent turns and straights of indeterminate direction; for example, the notation \mathbf{ts} means $\curvearrowright\uparrow$, $\curvearrowright\downarrow$, $\curvearrowleft\uparrow$, or $\curvearrowleft\downarrow$.

3 Controllability and existence

This section proves controllability of the differential drive, and existence of an optimal trajectory between any pair of start and goal configurations.

Theorem 1 *The bounded velocity differential drive in the unobstructed plane is controllable; i.e. there is a trajectory between any given pair of start and goal states.*

Proof: A trajectory of the form $\curvearrowright\uparrow\curvearrowright$ can reach any goal from any start configuration. The robot turns right until it is aligned with the goal position; it drives forward to the goal position; it turns right to the goal orientation. \square

Theorem 2 *For any given start and goal configuration of a bounded velocity differential drive in the unobstructed plane, there is an absolutely continuous time optimal trajectory and a corresponding measurable control.*

Proof: Theorem 6 of Sussmann and Tang [?] gives conditions sufficient for the existence of a time optimal control from a starting configuration s to a goal configuration g . For our case the conditions are:

- an admissible control exists taking the robot from s to g ;
- the system state variable $q = (x, y, \theta)$ takes values in an open subset of a differentiable manifold;
- the vector fields f_1 and f_2 are locally Lipschitz;
- the control $u = (\omega_1, \omega_2)$ takes values in a compact convex subset of \mathbf{R}^2 ;
- the admissible controls are measurable functions on compact subintervals of \mathbf{R} ;
- for every start state and every control over some time interval, there is a trajectory starting at the start state, and defined over the whole interval.

Theorem 1 shows that trajectories exist for every pair of given start and goal states, and the remaining conditions are readily verified. It follows that a time optimal trajectory and corresponding control exist for every given start and goal state. \square

4 Pontryagin's maximum principle

This section uses Pontryagin's maximum principle [?] to derive necessary conditions for time optimal trajectories of the bounded velocity differential drive robot. Section 5 provides a more geometric description of the optimal trajectories. This section logically precedes section 5, but perhaps it is easier to read them in parallel.

The maximum principle states that if the trajectory $q(t)$ with corresponding control $u(t)$ is time optimal then the following conditions must hold:

- There exists an *adjoint function*: an absolutely continuous \mathbf{R}^3 -valued function of time:

$$\lambda(t) = \begin{pmatrix} \lambda_1(t) \\ \lambda_2(t) \\ \lambda_3(t) \end{pmatrix}$$

- The adjoint function is nontrivial, *i.e.* not identically zero.
- The adjoint function satisfies the *adjoint equation*:

$$\dot{\lambda} = -\frac{\partial}{\partial q} H \quad \text{a.e.} \quad (15)$$

where $H : \mathbf{R}^3 \times SE^2 \times U \rightarrow \mathbf{R}$ is the *Hamiltonian*:

$$H(\lambda, q, u) = \langle \lambda, \dot{q}(q, u) \rangle \quad (16)$$

- The control $u(t)$ minimizes the Hamiltonian at almost every t :

$$H(\lambda(t), q(t), u(t)) = \min_{z \in U} H(\lambda(t), q(t), z) \quad \text{a.e.} \quad (17)$$

Equation 17 is called the *minimization equation*.

- We define $\lambda_0(t)$ to be the negative of the minimum attained for the Hamiltonian. So

$$\lambda_0(t) = -H(\lambda(t), q(t), u(t)) \quad \text{a.e.} \quad (18)$$

and λ_0 is constant in time and non-negative.

$$\lambda_0 \geq 0 \quad (19)$$

4.1 Application of the maximum principle to differential drives

To apply the maximum principle to the differential drive, we begin with the adjoint equation (15). We substitute the system equations (11) and the definition of the Hamiltonian (16) into the adjoint equation. After simplification we obtain:

$$\dot{\lambda} = -\frac{\partial}{\partial q}H = (0, 0, \lambda_1\dot{y} - \lambda_2\dot{x}) \quad \text{a.e.} \quad (20)$$

We can then integrate to solve for λ :

$$\lambda = (c_1, c_2, \eta(x, y)) \quad (21)$$

where (c_1, c_2, c_3) are arbitrary integration constants, and we define η to be

$$\eta(x, y) = c_1y - c_2x + c_3. \quad (22)$$

Since the maximum principle says the adjoint function is not identically zero, the integration constants c_1 , c_2 , and c_3 are not all zero, so in fact the adjoint function is never zero.

Having solved for the adjoint function, we now use the minimization equation (17) to derive the optimal controls. First we substitute the system equations (11) into the definition of the Hamiltonian (16):

$$H(\lambda, q, u) = \langle \lambda, \omega_1 f_1(q) + \omega_2 f_2(q) \rangle \quad (23)$$

$$= \omega_1 \langle \lambda, f_1(q) \rangle + \omega_2 \langle \lambda, f_2(q) \rangle \quad (24)$$

We define two *switching functions* φ_1 and φ_2 and an associated function φ_3 :

$$\varphi_1(t) = \langle \lambda(t), f_1(q(t)) \rangle \quad (25)$$

$$\varphi_2(t) = \langle \lambda(t), f_2(q(t)) \rangle \quad (26)$$

$$\varphi_3(t) = \langle \lambda(t), f_3(q(t)) \rangle \quad (27)$$

Substituting φ_1 and φ_2 into equation 24 yields:

$$H(\lambda(t), q(t), u(t)) = \omega_1(t)\varphi_1(t) + \omega_2(t)\varphi_2(t) \quad (28)$$

which is minimized by setting each wheel angular velocity ω_i to the negative sign of the corresponding switching function φ_i . (We define the function $\text{sign}(\cdot)$ to be multivalued at zero, taking on any value in the interval $[-1, 1]$.)

Theorem 3 *If $q(t)$ is a time optimal trajectory with control $u(t)$ for a bounded velocity differential drive robot, then there is a nontrivial absolutely continuous adjoint function $\lambda(t)$ satisfying the adjoint equation, and the control must satisfy*

$$\omega_1(t) = -\text{sign}(\varphi_1(t)) \quad (29)$$

$$\omega_2(t) = -\text{sign}(\varphi_2(t)) \quad (30)$$

for almost all t .

Proof: Application of the maximum principle. \square

An *extremal* is any trajectory and associated control that satisfies the maximum principle. Since the maximum principle is a necessary condition for time optimal trajectories, it follows that all time optimal trajectories are extremals, but not all extremals are time optimal. For example, we shall see that a turn in place through an angle of 2π is extremal, but it is obviously not time optimal.

Before focusing on time optimal trajectories, we develop more properties of the broader class of extremals. First, we observe that the three vectors $f_1(q)$, $f_2(q)$, and $f_3(q)$ are linearly independent for every q . So equations 25 through 27 define the φ_i to be the coordinates of λ using the f_i as a vector basis. We noted earlier that the adjoint function $\lambda(t)$ is never zero. Consequently we observe that

$$|\varphi_1| + |\varphi_2| + |\varphi_3| \neq 0 \quad (31)$$

for all t .

We also observe that if we substitute equations 29 and 30 into equation 28 we obtain

$$-H = |\varphi_1| + |\varphi_2| \quad \text{a.e.} \quad (32)$$

which is also equal to λ_0 for almost all t .

$$|\varphi_1| + |\varphi_2| = \lambda_0 \geq 0 \quad (33)$$

Equation 33 is true for all t , not just almost everywhere. The reason is that the $\varphi_i(\cdot)$ are absolutely continuous, since they are inner products of functions known to be absolutely continuous (equations 25 through 27).

4.2 The switching functions

The nature of the extremal controls is best explored by studying the switching functions φ_i . First we develop expressions for the time derivatives $\dot{\varphi}_i$, using a result from Sussmann and Tang [?]:

Lemma 1 (Sussmann and Tang’s Lemma 1)

Let f_Z be a smooth vector field and let $q(t)$ be a time optimal trajectory with associated control $u(t)$. Let φ_Z be given by

$$\varphi_Z(t) = \langle \lambda(t), f_Z(q(t)) \rangle \quad (34)$$

Then the time derivative is given by

$$\dot{\varphi}_Z = \omega_1 \langle \lambda, [f_1, f_Z] \rangle + \omega_2 \langle \lambda, [f_2, f_Z] \rangle \quad (35)$$

for almost all t .

Proof: given in [?]. \square

By substituting f_1 , f_2 , and f_3 for f_Z , we obtain expressions for $\dot{\varphi}_1$, $\dot{\varphi}_2$, and $\dot{\varphi}_3$, respectively:

$$\dot{\varphi}_1 = -\omega_2 \varphi_3 \quad (36)$$

$$\dot{\varphi}_2 = \omega_1 \varphi_3 \quad (37)$$

$$\dot{\varphi}_3 = \frac{1}{4b^2}(-\omega_1 + \omega_2)(\varphi_1 + \varphi_2) \quad (38)$$

for almost all t .

We need one other result before proceeding. By substituting the definitions of the φ_i , it is easily confirmed that for a time optimal trajectory

$$(\varphi_1 + \varphi_2)^2 + (2b\varphi_3)^2 = c_1^2 + c_2^2 \quad (39)$$

which is constant for the duration of the trajectory. We can subtract the square of equation 33 to obtain

$$2\varphi_1\varphi_2(1 - \text{sign}(\varphi_1\varphi_2)) + (2b\varphi_3)^2 = c_1^2 + c_2^2 - \lambda_0^2 \quad (40)$$

which is also constant for the duration of the trajectory.

We shall see that optimal trajectories are a sequence of *singular* and *generic* intervals. First we develop the singular intervals. We will say that an extremal trajectory is ω_1 -*singular* on some interval if $\varphi_1 \equiv 0$ on that interval. Similarly we will say that

an extremal trajectory is ω_2 -*singular* on some interval if $\varphi_2 \equiv 0$ on that interval. Finally, a trajectory is *doubly singular* on some interval when it is both ω_1 -singular and ω_2 -singular on that interval.

We will say that an extremal is *trivial* if it is motionless on its entire domain $[0, T]$. Note that if a trivial extremal is time optimal then its domain is the single instant $t = 0$.

Theorem 4 *If the switching functions φ_1 and φ_2 have a common root on an extremal trajectory of the bounded velocity differential drive, the trajectory is trivial.*

Proof: Let $q(t)$ be an extremal trajectory defined on the interval $[0, T]$, with $\varphi_1 = \varphi_2 = 0$ at some point. From equation 33, $\lambda_0 = 0$. Since λ_0 is invariant, it follows that φ_1 and φ_2 are identically zero on $[0, T]$. From equation 31 it follows that φ_3 is nonzero everywhere on $[0, T]$.

Since φ_1 is identically zero, $\dot{\varphi}_1$ is also identically zero. Since φ_3 is nonzero, equation 36 implies that ω_1 is zero for almost all t . Similar reasoning shows that ω_2 is zero for almost all t . So the robot never moves. \square

Theorem 5 *On a singular interval of an extremal of the bounded velocity differential drive, the controls ω_1 and ω_2 are constant and equal almost everywhere.*

Proof: If there is a common root of φ_1 and φ_2 , then by the previous theorem the entire trajectory is doubly singular and the result follows easily. We now consider the case of an ω_1 -singular interval with no common root of φ_1 and φ_2 . φ_1 is identically zero, so from equation 33 it follows that

$$|\varphi_2| = \lambda_0 > 0 \quad (41)$$

so φ_2 is constant. Likewise, since $\omega_2 = -\text{sign}(\varphi_2)$ (equation 30) almost everywhere on the interval, ω_2 is constant, either 1 or -1 almost everywhere on the singular interval.

Now we note that since φ_1 is identically zero, so is $\dot{\varphi}_1$. From equation 36 and the fact that ω_2 is nonzero almost everywhere, we conclude that φ_3 is identically zero, and so is its time derivative:

$$\varphi_3 \equiv 0, \dot{\varphi}_3 \equiv 0 \quad (42)$$

(Equation 42 is true for all t , not just almost everywhere, because φ_3 is absolutely continuous.) From equation 38 and the fact that $(\varphi_1 + \varphi_2)$ is nonzero, we conclude that $(-\omega_1 + \omega_2)$ is zero for almost all t , or

$$\omega_1 \equiv \omega_2 \quad \text{a.e.} \quad (43)$$

That completes the proof for an ω_1 -singular interval. The proof for an ω_2 -singular interval is similar. \square

Theorems 4 and 5 have shown that the controls on singular intervals are simple. We now turn our attention away from singular intervals, to show that the rest of an extremal is simple as well. We will use the term *generic point* to refer to any time that is not a root of either φ_1 or φ_2 . We will use the term *generic interval* to refer to an interval on which every time is a generic point.

Theorem 6 *For an extremal trajectory and corresponding control of the bounded velocity differential drive, there is a $\delta > 0$, such that every generic point is in a generic interval of duration at least δ . On a generic interval the controls are constant.*

Proof: The roots of $\varphi_1\varphi_2$ are roots of either φ_1 or φ_2 , so generic points are those points that are not roots of $\varphi_1\varphi_2$. The generic points are given by the set $(\varphi_1\varphi_2)^{-1}(\mathbf{R}^+ \cup \mathbf{R}^-)$. Because $(\mathbf{R}^+ \cup \mathbf{R}^-)$ is open and $\varphi_1\varphi_2$ is continuous, the set of generic points is also open, relative to the time interval on which the trajectory is defined. So the generic points are a union of open intervals, except possibly the first interval and last interval which are only half open. Consider an arbitrary generic point, and let I be the corresponding interval. The switching functions φ_1 and φ_2 do not change signs on I , and the controls ω_1 and ω_2 are constant almost everywhere on I .

Case 1: open interval, φ_1 and φ_2 have the same sign. It follows that $\omega_1 = \omega_2$ almost everywhere, so from equation 38 we know that $\dot{\varphi}_3$ is zero almost everywhere, so φ_3 is constant. From equations 36 and 37 we obtain

$$\dot{\varphi}_1 = -\dot{\varphi}_2 = \pm\varphi_3 \quad \text{a.e.} \quad (44)$$

If φ_3 is zero, then all the switching functions are constant. The interval cannot be bounded by roots of

φ_1 and φ_2 , so the extremal is just one big generic interval.

If φ_3 is nonzero, then the interval I is bounded by roots of either φ_1 or φ_2 , and each is linear in t . Further we know that

$$|\varphi_1 + \varphi_2| = \lambda_0 \quad (45)$$

So at one end of the interval we have $\varphi_1 = 0$, $\varphi_2 = \pm\lambda_0$, and at the other end of the interval we have $\varphi_1 = \pm\lambda_0$, $\varphi_2 = 0$. The length of the interval is defined to be δ_1 :

$$\delta_1 = |\lambda_0/\varphi_3| \quad (46)$$

Having excluded the case of $\varphi_3 = 0$, we can solve equation 40 for φ_3 and substitute:

$$\delta_1 = \frac{2b\lambda_0}{\sqrt{c_1^2 + c_2^2 - \lambda_0^2}} \quad (47)$$

Case 2: open interval, φ_1 and φ_2 have opposite signs. Let s be the sign of φ_1 , either 1 or -1 . Substituting s for ω_1 and $-s$ for ω_2 in equations 36 through 38 we obtain:

$$\dot{\varphi}_1 = s\varphi_3 \quad (48)$$

$$\dot{\varphi}_2 = s\varphi_3 \quad (49)$$

$$\dot{\varphi}_3 = -\frac{s}{2b^2}(\varphi_1 + \varphi_2) \quad (50)$$

Differentiating equation 50 gives

$$\ddot{\varphi}_3 = -\frac{s}{2b^2}(\dot{\varphi}_1 + \dot{\varphi}_2) \quad (51)$$

Substituting from equations 48 and 49 gives

$$\ddot{\varphi}_3 = -\frac{s}{2b^2}(2s\varphi_3) = -\frac{1}{b^2}\varphi_3 \quad (52)$$

since $s^2 = 1$. The solution is of the form

$$\varphi_3 = A \cos\left(\frac{t}{b} + t_n\right) \quad (53)$$

with constants A and t_n to be determined. Differentiating gives

$$\dot{\varphi}_3 = -\frac{A}{b} \sin\left(\frac{t}{b} + t_n\right) \quad (54)$$

Comparing with equation 50 gives

$$-\frac{A}{b} \sin\left(\frac{t}{b} + t_n\right) = -\frac{s}{2b^2}(\varphi_1 + \varphi_2) \quad (55)$$

Solving for $\varphi_1 + \varphi_2$:

$$\varphi_1 + \varphi_2 = 2Abs \sin\left(\frac{t}{b} + t_n\right) \quad (56)$$

Substituting into equation 40 and solving for A :

$$A = \frac{\sqrt{c_1^2 + c_2^2}}{2b} \quad (57)$$

So

$$\varphi_1 + \varphi_2 = s\sqrt{c_1^2 + c_2^2} \sin\left(\frac{t}{b} + t_n\right) \quad (58)$$

Now we rewrite equation 33 as

$$\varphi_1 - \varphi_2 = -s\lambda_0 \quad (59)$$

Adding and subtracting equations 58 and 59, we can solve for φ_1 and φ_2 , respectively:

$$\varphi_1 = s \left(\frac{\sqrt{c_1^2 + c_2^2}}{2} \sin\left(\frac{t}{b} + t_n\right) - \frac{\lambda_0}{2} \right) \quad (60)$$

$$\varphi_2 = s \left(\frac{\sqrt{c_1^2 + c_2^2}}{2} \sin\left(\frac{t}{b} + t_n\right) + \frac{\lambda_0}{2} \right) \quad (61)$$

Roots of φ_1 or φ_2 correspond to solutions of

$$\sin\left(\frac{t}{b} + t_n\right) = \pm \frac{\lambda_0}{\sqrt{c_1^2 + c_2^2}} \quad (62)$$

The present case assumes an open interval bounded by roots of φ_1 or φ_2 , which implies that $\lambda_0^2 \leq c_1^2 + c_2^2$. If $\lambda_0^2 < c_1^2 + c_2^2$ the solutions are separated either by δ_2 or by $\pi - \delta_2$, where we define δ_2 as

$$\delta_2 = 2b \sin^{-1} \left(\frac{\lambda_0}{\sqrt{c_1^2 + c_2^2}} \right) \quad (63)$$

On intervals of length $\pi - \delta_2$, φ_1 and φ_2 have the same sign, so for the present case the interval is of length δ_2 . If $\lambda_0^2 = c_1^2 + c_2^2$ the solutions are separated by $b\pi$, a case already covered by the definition of δ_2 . So if the interval I is open, with φ_1 and φ_2 having the same sign, the length of I is δ_2 .

Other cases: closed and half closed intervals. The other cases are readily dealt with. If the generic interval is closed then it must be the entire extremal, and the theorem is trivially true. The remaining cases are half open intervals, which might be present at the beginning or end of the trajectory. The proof is completed by defining δ to be the minimum of δ_1 , δ_2 , and the length of the half-open intervals if present. \square

The last element we need to describe extremals is the *switching*. By a *switching* we mean a point such that for no neighborhood of the point are the controls constant almost everywhere on the neighborhood. Theorems 4 through 6 show that the controls on singular and generic intervals are constant. It follows that every switching is a root of a switching function, and is not interior to a singular interval (or a generic interval).

We need to show that extremals have only a finite number of switchings. If extremals were made up of generic intervals only, then the lower bound on the length of generic intervals would imply that every extremal has a finite number of switchings. Unfortunately singular intervals are not as simple: there is no lower bound on their length. We must also exclude the possibility that two switchings might be separated by neither a generic nor a singular interval. The following lemma will enable us to deal with both problems.

Lemma 2 *For every extremal trajectory of the bounded velocity differential drive, if two points are given, each of them a root of a switching function, the given points either share a common singular interval or have a generic point between them.*

Proof: Let t_1 and t_2 be the given points, and suppose there are no generic points between them. Consider the extremal obtained by restricting the given trajectory to the interval $[t_1, t_2]$. Every point on the restricted extremal is a root of one of the switching functions. Let $R_1 = \varphi_1^{-1}(0)$ be the roots of the first switching function, and let $R_2 = \varphi_2^{-1}(0)$ be the roots of the second switching function. R_1 and R_2 are closed and their union covers the interval. If the

intersection $R_1 \cap R_2$ is nonempty, points in that intersection are roots of both switching functions, so that the entire trajectory is doubly singular. If the intersection $R_1 \cap R_2$ is empty, then R_1 and R_2 are complements. Each is both open and closed, relative to the interval $[t_1, t_2]$. One must be empty, and the other must be the entire interval, so that the entire interval $[t_1, t_2]$ is singular, and the two points share a common singular interval. \square

Theorem 7 *For an extremal trajectory defined on the interval $[0, T]$, the interval is covered by the closures of a finite number of generic or singular intervals. There are a finite number of switchings, occurring only at the ends of the intervals. Every interval is one of three types:*

- Doubly singular. *Both switching functions are identically zero. Both controls are zero for almost all time—the robot never moves. An extremal with a double singularity consists of a single interval—it is doubly singular throughout.*
- Singular. *One of the switching functions φ_1 or φ_2 is identically zero, and the other is never zero. The controls are equal and saturated (1 or -1) almost everywhere, so the robot goes either straight forward (\uparrow) or straight backward (\downarrow) at full speed.*
- Generic. *Neither switching function has a root anywhere on a generic interval. There are two cases:*
 - *The switching functions have the same sign. The controls are equal and saturated almost everywhere, so, as with the singular interval, the robot goes either straight forward (\uparrow) or straight backward (\downarrow) at full speed.*
 - *The switching functions have opposite sign. The controls are opposite and saturated almost everywhere, so the robot turns in place, either counterclockwise (\curvearrowright) or clockwise (\curvearrowleft).*

Proof: First we show that the interval is covered by the closures of the generic and singular intervals.

Suppose otherwise. Then the difference is nonempty, and is relatively open, comprising a union of intervals open relative to $[0, T]$. Take two points in one of these open intervals. By lemma 2 they are separated by a generic point, a contradiction. We conclude that the closures of the generic and singular intervals comprise the entire trajectory.

Define the set of *delimiting points* to include every root of a switching function that is not interior to a singular interval. We now show that the delimiting points are finite in number. By theorem 6 we know there is a lower bound δ on the length of generic intervals. Suppose we are given three delimiting points. The first and third must be separated by at least δ , otherwise the second would be in the interior of a singular interval by lemma 2. It follows there are no more than $(2T/\delta) + 1$ delimiting points.

The delimiting points, with the addition of the endpoints 0 and T , define the structure of the extremal. They delimit a finite sequence of intervals. Each interval is easily shown to be either extremal or generic, and each is maximal, being contained in no larger generic or extremal interval. Since switchings occur only at delimiting points, the switchings are also finite in number.

The rest follows easily. By theorems 4 and 5 the controls are constant on singular intervals, and by theorem 6 the controls are constant on generic intervals. The rest of the theorem follows from the case structure introduced in the proofs of theorems 4, 5, and 6. \square

We note that although every switching is a delimiting point, not every delimiting point is a switching. There are extremal trajectories, even time optimal trajectories, comprising two generic intervals separated by a delimiting point that is not a switching.

5 Geometric interpretation of extremals

The previous section defined extremal trajectories, a superset of time optimal trajectories, and showed that every nontrivial extremal trajectory is a finite sequence of straights and turns in place. This section

reexamines the structure of extremals with a geometrical perspective.

The central geometrical construct is a directed line in the plane we will call the η -line. Recall (equation 22) that $\eta(x, y) = c_1y - c_2x + c_3$, where c_1 , c_2 , and c_3 are three real constants whose existence is given by the maximum principle, and which cannot all be zero. Assuming that c_1 and c_2 are not both zero, we define the η -line to be the locus of points (x, y) such that $\eta(x, y)$ is zero. Then for all (x, y) in the plane we can interpret $\eta(x, y)$ as the signed distance of the point from the η -line, scaled by $\sqrt{c_1^2 + c_2^2}$. We choose the direction of the η -line so that positive values of η are to the right, and negative values of η are to the left.

We can also include $c_1 = c_2 = 0$ as a special case. In that case we will say the η -line is at infinity, and the entire plane is to the right of the η -line if c_3 is positive, or to the left of the η -line if c_3 is negative.

Now it turns out that for an extremal trajectory the control laws (equations 29 and 30) can be interpreted geometrically:

$$\omega_1 \begin{cases} = 1 & \text{if wheel 2} \in \text{right half plane} \\ \in [-1, 1] & \text{if wheel 2} \in \eta\text{-line} \\ = -1 & \text{if wheel 2} \in \text{left half plane} \end{cases} \quad (64)$$

$$\omega_2 \begin{cases} = 1 & \text{if wheel 1} \in \text{left half plane} \\ \in [-1, 1] & \text{if wheel 1} \in \eta\text{-line} \\ = -1 & \text{if wheel 1} \in \text{right half plane} \end{cases} \quad (65)$$

The left wheel switches only when the right wheel touches the η -line, and the right wheel switches only when the left wheel touches the η -line.

The proof is as follows. Recall that the controls are given by the signs of two *switching functions* $\varphi_1(t)$ and $\varphi_2(t)$. From equations 29 and 30:

$$\omega_1(t) = -\text{sign}(\varphi_1(t))$$

$$\omega_2(t) = -\text{sign}(\varphi_2(t))$$

Define the wheel coordinates (x_1, y_1) for wheel 1 and (x_2, y_2) for wheel 2. Specifically,

$$\begin{pmatrix} x_1 \\ y_1 \end{pmatrix} = \begin{pmatrix} x - b \sin \theta \\ y + b \cos \theta \end{pmatrix} \quad (66)$$

$$\begin{pmatrix} x_2 \\ y_2 \end{pmatrix} = \begin{pmatrix} x + b \sin \theta \\ y - b \cos \theta \end{pmatrix} \quad (67)$$




Figure 3: The switching functions φ_1 and φ_2 give the scaled signed distances of wheels 1 and 2 from the η -line.

Then the switching functions can be written

$$\varphi_1 = -\frac{1}{2b}\eta(x_2, y_2) \quad (68)$$

$$\varphi_2 = \frac{1}{2b}\eta(x_1, y_1) \quad (69)$$

which is confirmed by expanding the definitions of the switching functions (equations 25 through 27), using the definitions of the vector fields (equation 10), the expression derived for the adjoint function (equation 21), the definition of $\eta(x, y)$ (equation 22), and also substituting equations 66 and 67 above.

Equations 68 and 69 have a geometrical interpretation illustrated in figure 3. φ_1 is the negative signed distance of wheel 2 from the η -line, scaled by $\sqrt{c_1^2 + c_2^2}/2b$. φ_2 is the positive signed distance of wheel 1 from the η -line, scaled by the same factor. Substituting into the control equations (29 and 30):

$$\omega_1 = \text{sign}(\eta(x_2, y_2)) \quad (70)$$

$$\omega_2 = -\text{sign}(\eta(x_1, y_1)) \quad (71)$$

where $\text{sign}(\cdot)$ may be set-valued, returning the entire interval $[-1, 1]$ as the value of $\text{sign}(0)$. The reader may readily confirm these are equivalent to equations 64 and 65, completing the proof.

The result is a geometric version of the main result of the previous section, theorem 7: For an extremal trajectory defined on an interval $[0, T]$, the interval is a sequence of intervals of the following types:

- *Doubly singular.* The robot sits motionless with both wheels exactly on the η -line. In this case the entire extremal is a single interval.
- *Singular.* The robot rolls straight and parallel to the η -line, either forward (\uparrow) or backward (\downarrow) with one wheel exactly on the η -line.
- *Generic*
 - The switching functions have the same sign. The robot straddles the η -line, and rolls either straight forward (\uparrow) or straight backward (\downarrow), whichever direction is closer to the η -line direction.
 - The switching functions have opposite sign. The wheels are on the same side of the η -line. The robot turns in place, either clockwise (\curvearrowright) if the robot is to the right of the η -line, or counterclockwise (\curvearrowleft) if the robot is to the left of the η -line.

Thus the robot’s motion depends on the location of the η -line, but the location of the η -line is not easily found. It is defined by the three integration constants c_1 , c_2 , and c_3 . The maximum principle tells us these constants exist, but not how to find them. Evidently there is no direct method for finding the location of the η -line. Its location is determined in the course of finding the time optimal trajectories, which later sections will address.

The geometrical interpretation of the switching functions yields some further insights into the analysis of the previous section. We have already seen that φ_1 and φ_2 measure the distances of the wheels from the η -line. There is also a geometric interpretation for φ_3 . Let β be the robot heading relative to the η -line.

$$\beta = \theta - \tan^{-1}(-c_2, c_1)$$

Then φ_3 is the scaled sine of β . It is zero when the robot’s heading is parallel to the η -line. Hence, as observed in the proof of theorem 6, when the robot rolls straight φ_3 is stationary while φ_1 and φ_2 vary linearly, and when the robot turns all of the φ_i vary sinusoidally.




Figure 4: Two extremals: *tangent CW* and *zigzag right*. Other extremal types are *zigzag left*, *tangent CCW*, and turning in place: *CW* and *CCW*. Straight lines are special cases of zigzags or tangents.

We also note that equation 40 has a geometric interpretation: it is Pythagoras’s theorem applied to a triangle whose hypotenuse is the robot wheelbase.

6 Enumeration of extremals

So far we have identified several different types of generic and singular intervals. However, not all sequences of intervals are valid. The robot switches controls only when a wheel touches the η -line. Therefore the extremals can be enumerated by constructing a circle whose diameter is the robot wheelbase. Depending on whether this circle is far from the η -line, tangent to the η -line, or crossing the η -line, different patterns arise:

- **CW and CCW:** If the wheelbase circle is entirely in the open right half plane then the robot turns in the clockwise direction and can never touch the η -line (case CW). There are no switchings. CCW is similar. The structure of CW trajectories is \curvearrowright ; the structure of CCW trajectories is \curvearrowleft .
- **TCW and TCCW (Tangent CW and Tangent CCW).** See figure 4. TCW arises if the wheel base circle is in the closed right half-plane and tangent to the η -line. If no wheel is touching the η -line

the robot turns clockwise. If a wheel is touching the η -line then the robot may either roll straight along the η -line, or it may turn clockwise. Throughout this motion, the wheelbase circle is always tangent to the η -line and in the closed right half-plane. The structure of TCW trajectories is an alternation of singular straights and clockwise turns (\curvearrowright). The turns are all integer multiples of $-\pi$, except possibly the first and last. TCCW is similar.

- **ZR and ZL (Zigzag right and zigzag left).** If the η -line intersects the interior of the wheel base circle then a zigzag behavior occurs, alternating straights and turns. If the robot initially straddles the η -line, it rolls straight ahead until one wheel touches the η -line. It then turns until the other wheel touches the η -line, and then goes straight again. There are two non-degenerate patterns: $\dots \uparrow \curvearrowright \downarrow \curvearrowleft \dots$ called *zigzag right (ZR)*, and $\dots \uparrow \curvearrowleft \downarrow \curvearrowright \dots$ called *zigzag left (ZL)*. Degenerate cases occur when both wheels are touching the line (the trivial doubly singular case) or when the robot heading is parallel to the η -line, in which case it goes straight and never switches.

Every extremal trajectory, and hence every time optimal trajectory, is one of the types above. For trajectories of just one or two actions it may not be possible to identify the type unambiguously. For example a straight line, no matter how long, could be TCW, TCCW, ZR, or ZL. But in most cases it is possible to tell the difference. We will see that most time optimal trajectories have three actions, which is enough to determine the type. If the turns are all in the same direction it is a tangent. If the turns alternate direction it is a zigzag. If the straights are all parallel it is a tangent. If the straights are not parallel it is a zigzag.

7 Comparison between steered cars and differential drives

Section 5 provided a geometric interpretation of differential drive extremals. Section 6 used the geometry to classify and enumerate the extremals. In this

section we generalize the geometric analysis and apply it to steered cars as well as differential drives.

Dubins [?], Reeds and Shepp [?], Sussman and Tang [?], and others have presented the time optimal trajectories for steered cars with bounded velocity and a constraint on the minimum turning radius. It might at first seem that a differential drive is just a steered car with a minimum turning radius of zero. Then we would expect the optimal differential drive trajectories to be limits of the optimal steered car trajectories as the minimum turning radius approaches zero.

However, this assumption quickly leads to contradictions. Turns in place would take zero time, and trajectories of the form “turn to face the goal, drive forwards, turn to the desired angle” would always be optimal. In fact, the robot could spin in place any number of times before driving towards the goal, and during any spin in place, the wheel speeds would be infinite. In order to illuminate the similarities and differences between the optimal trajectories for bounded velocity differential drives and steered cars, we will first consider a common choice of controls.

Velocity and angular velocity controls

For steered cars, the controls are velocity and angular velocity. We may choose the controls for the differential drive to be (v, ω) also. For both systems,

$$\dot{q} = \begin{pmatrix} v \cos \theta \\ v \sin \theta \\ \omega \end{pmatrix} \quad (72)$$

Reeds and Shepp assumed the control bounds

$$v = \pm 1 \quad |\omega| \leq \alpha^2 \quad (73)$$

The parameter α^2 expresses the constraint on turning radius. It is sufficient for our purposes to assume that $\alpha = 1$. The admissible controls fall on two line segments in (v, ω) space. (See figure 5.)

For the differential drive, the wheel speeds are bounded: $(\omega_1, \omega_2) \in U$. We use equations 4 and 5 to map U to (v, ω) space. The result is a diamond, shown in figure 5. The corners of U map to the vertices of the diamond. As we might expect, the max-




Figure 5: Control spaces for the Dubins car, the Reeds and Shepp car, and the differential drive.

imum angular velocity ω is inversely proportional to the wheelbase of the robot b .

7.1 The maximum principle

We now apply the maximum principle to equation 72, which holds for both the steered car and the differential drive. The Hamiltonian for the system is the inner product of the adjoint vector λ with the state derivative \dot{q}

$$H = \langle \lambda, \dot{q} \rangle = \lambda_1 v \cos \theta + \lambda_2 v \sin \theta + \lambda_3 \omega \quad (74)$$

We write the adjoint equation

$$\dot{\lambda} = -\frac{\partial}{\partial \mathbf{q}} H = (0, 0, \lambda_1 v \sin \theta - \lambda_2 v \cos \theta) \quad (75)$$

and integrate to solve for λ

$$\lambda = (c_1, c_2, c_1 y - c_2 x + c_3) \quad (76)$$

Substituting equation 76 into the Hamiltonian,

$$H = v(c_1 \cos \theta + c_2 \sin \theta) + \omega(c_1 y - c_2 x + c_3) \quad (77)$$

Notice that $c_1 y - c_2 x + c_3 = \eta(x, y)$. For simplicity, we will assume that

$$c_1^2 + c_2^2 = 1 \quad (78)$$

Then $\eta(x, y)$ gives the signed distance of the reference point of the robot from the η -line. Recall that β is the orientation of the robot relative to the η -line:

$$\beta = \theta - \text{atan}(-c_2, c_1) \quad (79)$$

Substituting into equation 77,

$$-H = v \cos \beta + \omega \eta \quad (80)$$

Equation 80 is similar to equation 28: both express the Hamiltonian as a dot product between the controls and some function of the configuration. For the steered car, it turns out that $\cos \beta$ and η are the switching functions.

7.2 Level sets of the Hamiltonian

The right hand side of equation 80 is a dot product between two vectors. Define

$$w = (v, \omega) \quad (81)$$

$$d = (\cos \beta, \eta) \quad (82)$$

w is the control, and we call d the *characteristic vector*. From the maximum principle,

$$\lambda_0 = \max(-H(t)) \quad (83)$$

$$\langle w, d \rangle = \lambda_0 \quad (84)$$

That is, the control w must minimize the dot product, and the dot product is constant.

Equations 83 and 84 lead to a useful geometric construction. We first consider the differential drive. Choose a small but nonzero value for λ_0 ($0 < \lambda_0 < 1$), and consider the point $w_1 = (1, 0)$ on the control boundary, corresponding to the robot driving in a straight line. There is a line of values for the characteristic vector that make a constant dot product with w_1 , labelled “line 1” in figure ???. Similarly, consider the other three vertices of the control boundary and construct three corresponding lines. If d falls along one of the lines, the dot product in 84 will be equal

to λ_0 for at least one of the controls w_1, w_2, w_3 , or w_4 .

The maximum principle gives an additional constraint, described by equation 83. The fact that the control must be maximizing restricts the possible values of the characteristic vector to the rectangle formed by the four lines, shown in figure ???. No new points are added to or removed from the rectangle when we consider the remaining admissible controls.

If we consider the Hamiltonian as a function of the characteristic vector

$$H(d) = -\langle w(d), d \rangle \quad (85)$$

then the constructed rectangle is a level set of the Hamiltonian:

$$d \in H^{-1}(-\lambda_0) \quad (86)$$

We may find other level sets of the Hamiltonian by scaling λ_0 . Geometrically, increasing λ_0 uniformly scales the constructed rectangle. However, there is one additional constraint: the first element of the characteristic vector d is $\cos \beta \leq 1$. If λ_0 is large, then parts of the constructed rectangle must be clipped to satisfy this constraint. The first box of figure 6 shows some level sets of the Hamiltonian for the differential drive.

In order to satisfy the maximum principle, the characteristic vector d must fall on a single level set of the Hamiltonian. Furthermore, the characteristic vector is an absolutely continuous function of time. If the level set of the Hamiltonian is comprised of disjoint sections (λ_0 large), then the characteristic vector will be restricted to one of the sections.

Since the level sets of the Hamiltonian are qualitatively different for large and small λ_0 , we expect the extremal trajectories to be qualitatively different for large and small λ_0 . In fact, the value of λ_0 provides a convenient way to distinguish the classes of trajectories CW, CCW, TCW, TCCW, ZL, and ZR described in the previous section.

- $\lambda_0 > 1$: The robot is far from the η -line, and the maximizing control will always be one of \curvearrowright or \curvearrowleft . This corresponds to the trajectory classes CW and CCW, and is shown in the first column, second row of figure 6.

- $\lambda_0 = 1$: One of the robot's wheels falls on the η -line. This corresponds to the trajectory classes TCW and TCCW, and is shown in the first column, third row of figure 6.

- $\lambda_0 < 1$: The robot is straddling the η -line. This corresponds to the trajectory classes ZL and ZR, and is shown in the first column, fourth row of figure 6.

We may apply the same geometric approach to the steered car. To find candidate values for the characteristic vector, we choose a value for λ_0 and consider each of the possible controls. For small λ_0 , the level sets of the Hamiltonian are diamonds in the control space. For large λ_0 , we clip the diamonds to satisfy the constraint that $|\cos \beta| \leq 1$. The first row, second column of figure 6 shows the result. Once the level sets are constructed, we may again classify the trajectories by the value of λ_0 :

- $\lambda_0 > 1$: The robot is far from the η -line, and the maximizing control will either always maximize or minimize ω . The robot follows successive half-circles, switching driving direction but maintaining a constant angular velocity. The second column, second row of figure 6 shows an example.

- $\lambda_0 = 1$: The robot is always either following a $\pi/2$ arc of a circle with an endpoint on the η -line, or driving in a straight line along the η -line. See the second column, third row of figure 6.

- $\lambda_0 < 1$: The robot is near the η -line, and makes a sharp angle with the η -line. The extremal trajectory resembles a parallel parking maneuver. See the second column, fourth row of figure 6.

7.3 The Hamiltonian as power

We provide one more interpretation of the maximum principle that applies to both the differential drive and the steered car. The right hand side of equation 80 is a dot product between a generalized velocity and vector dependent on the location of the robot relative to some line. Consider a line of force coincident with the η -line, but with opposite direction.




Figure 6: Classification of extremals by λ_0 and level sets of the Hamiltonian, for both the differential drive (left column) and the steered car (right column.)

Assume a force of 1 is applied against the robot, pushing on the robot in the direction opposite the η -line, and causing a torque about the center of the robot. We could find the work done by the robot against this force, if we were given a trajectory for the robot. The time derivative of the work is the power. With a suitable choice of units,

$$p = (v \cos \beta + \omega \eta) \quad (87)$$

The right hand side of this equation and the right hand side of equation 80 are the same.

Informally, the maximum principle then tells us that the time optimal trajectories for both the differential drive and the steered car maximize power along or against some line of force in the plane; the η -line is this line of force.

8 Further conditions on time optimal trajectories

This section turns to the paper's primary focus: time optimal trajectories of bounded velocity differential drive vehicles. Every time optimal trajectory is an extremal, and therefore must be of type CCW, CW, TCCW, TCW, ZR, or ZL. In this section we find additional necessary conditions optimal trajectories must satisfy, ultimately finding that no time optimal trajectory can have more than three straights and two turns. We use these results to enumerate optimal trajectory types.

We first make a useful observation about optimal trajectories. Recall equations 13 and 14: $s(t)$ is the rectified arc length of the robot's path measured in the plane \mathbf{E}^2 , and $\sigma(t)$ is the rectified arc length of the robot's orientation measured in \mathbf{S}^1 .

Theorem 8 *For an optimal trajectory defined on the interval $[0, T]$, for all $t \in [0, T]$,*

$$t = s(t) + b\sigma(t) \quad (88)$$

Proof: By theorem 7 a nontrivial extremal trajectory is composed of intervals which are either turning in place or going straight at maximum velocity. The time required for each straight is just the length, and




Figure 7: There is always a roundtrip of length 2π visiting θ_s , θ_g , 0, and π . That gives two paths from θ_s , through 0 or π , to θ_g . At least one of them has length no more than π .

the time for each turn is b times the arc length. The result applies immediately for nontrivial time optimal trajectories, and also applies to a trivial (doubly singular) time optimal trajectory since in that case T is zero. \square

In [?], we demonstrate the more general result that equation 88 actually holds for any trajectory such that $\max(|\omega_1(t)|, |\omega_2(t)|) = 1$ almost everywhere; *i.e.*, for trajectories in which one control is always saturated. This may provide some intuition for why spins in place and straight lines are faster than curves. The next result gives a limit on how much time a robot following an optimal trajectory may spend spinning in place.

Theorem 9 *For every time optimal trajectory $\sigma(T) \leq \pi$.*

Proof: Consider the fastest **tst** trajectory between a given start and goal. Obviously the straight action connects the start to the goal. Suppose we orient our coordinates so the angle from start to goal is 0. The robot's heading during the straight is either 0 or π . Consider the robot's heading over the course of the entire trajectory. It starts at θ_s , passes through either 0 or π , and ends at θ_g . The fastest **tst** trajectory minimizes the arc length of a path on the circle that satisfies those constraints. We can show that the total arc length never exceeds π . Consider the possible paths on the circle illustrated in figure 7. Note that in every instance there is a round trip from θ_s to 0, π , θ_g , and back to θ_s , of length 2π . Hence there is a one-




Figure 8: Zigzags of three turns are not optimal

way trip from θ_s , through 0 or π , to θ_g of length π or less. This completes the proof; the **tst** trajectory we have constructed will be shorter than any trajectory with $\sigma(T) > \pi$. \square

Theorem 10 *Tangent trajectories containing more than three actions are not optimal.*

Proof: An extremal of type TCW or TCCW alternates turns and straights. Any full untruncated turn must be a multiple of π . If there are four actions, there is at least one untruncated turn of length at least π , and a second turn of nonzero length. The rectified arc length σ would be more than π , contradicting theorem 9. \square

Zigzag trajectories are composed of alternating turns and straights. Successive turns or straights must be in opposite directions, but have the same magnitude if untruncated. Simple geometry also gives a relationship between ϕ , the angle of each turn, and d , the length of each straight:

$$d = 2b \tan\left(\frac{\phi}{2}\right) \quad (89)$$

Theorem 11 *Zigzag trajectories containing three turns are not optimal.*

Proof: Consider a zigzag with three turns, and two straights. The straights are the same length, so the path comprises two legs of an isosceles triangle. Construct the circle containing the start, the goal, and the via point as in figure 8. If we perturb the via point to a nearby point on the same circle the turning time is unchanged, and the translation is decreased. Hence the original zigzag was not optimal. \square




Figure 9: Periodicity of a zigzag

Zigzags are also periodic in robot heading and distance from the η -line. Recall that β is the robot heading relative to the η -line and η is the signed scaled distance from the η -line. Let τ be the smallest positive time such that:

$$\begin{aligned} \beta(t) &= \beta(t + \tau) \\ \eta(t) &= \eta(t + \tau) \end{aligned}$$

Theorem 12 *A zigzag trajectory of more than one period is not optimal.*

Proof: Consider a zigzag of more than one period, beginning at time 0 and ending at time $T > \tau$. By theorem 11, the zigzag is not optimal if $\sigma(T) > 2\phi$. If $s(T) > 2d$, then there are three straights. The first and last straights are parallel. If we reorder the actions to perform these consecutively, we get a straight with length greater than d . Thus we have a trajectory which costs no more than the original but which is no longer a legitimate zigzag, nor any other extremal trajectory. Since it is not extremal, neither it nor the original trajectory can be optimal. \square

We apply theorems 10, 11, and 12 to the extremal trajectory types CCW, CW, TCCW, TCW, ZR, or ZL to enumerate the types of optimal trajectories. Table 1 shows the results.

9 Symmetries

Further analysis of the time optimal trajectories is difficult because of the large number of cases. This

	base	T_1	T_2	T_3	$T_2 \circ T_1$	$T_3 \circ T_1$	$T_3 \circ T_2$	$T_3 \circ T_2 \circ T_1$
A.	\curvearrowright	\curvearrowright	\curvearrowright	\curvearrowright	\curvearrowright	\curvearrowright	\curvearrowright	\curvearrowright
B.	\downarrow	\uparrow	\downarrow	\downarrow	\uparrow	\uparrow	\downarrow	\uparrow
C.	$\downarrow\curvearrowright$	$\uparrow\curvearrowright$	$\curvearrowright\downarrow$	$\downarrow\curvearrowright$	$\curvearrowright\uparrow$	$\uparrow\curvearrowright$	$\curvearrowright\downarrow$	$\curvearrowright\uparrow$
D.	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$
E.	$\uparrow\curvearrowright\pi\downarrow$	$\downarrow\curvearrowright\pi\uparrow$	$\downarrow\curvearrowright\pi\uparrow$	$\uparrow\curvearrowright\pi\downarrow$	$\uparrow\curvearrowright\pi\downarrow$	$\downarrow\curvearrowright\pi\uparrow$	$\downarrow\curvearrowright\pi\uparrow$	$\uparrow\curvearrowright\pi\downarrow$
F.	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$
G.	$\downarrow\curvearrowright\uparrow$	$\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow$	$\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow$	$\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow$	$\downarrow\curvearrowright\uparrow$
H.	$\curvearrowright\downarrow\curvearrowright\uparrow$	$\curvearrowright\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow\curvearrowright$	$\downarrow\curvearrowright\uparrow\curvearrowright$	$\downarrow\curvearrowright\uparrow\curvearrowright$	$\uparrow\curvearrowright\downarrow\curvearrowright$	$\uparrow\curvearrowright\downarrow\curvearrowright$	$\downarrow\curvearrowright\uparrow\curvearrowright$
I.	$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$

Table 2: The forty optimal trajectory types sorted by symmetry class.

$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$	$\curvearrowright\uparrow\curvearrowright$	$\curvearrowright\downarrow\curvearrowright$
$\uparrow\curvearrowright\pi\downarrow$	$\downarrow\curvearrowright\pi\uparrow$	$\uparrow\curvearrowright\pi\downarrow$	$\downarrow\curvearrowright\pi\uparrow$
$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$	$\uparrow\curvearrowright\downarrow\curvearrowright\uparrow$	$\downarrow\curvearrowright\uparrow\curvearrowright\downarrow$

Table 1: Enumeration of optimal trajectory types. The structure of any optimal trajectory must be one of the above, or a subsection of one of the above.

complexity is reduced using symmetries developed by Souères and Boissonnat [?] and Souères and Laumond [?] for steered cars.

The symmetries are summarized in figure 10. Let u be a control from $q = (x, y, \theta)$ to the origin. Then there are seven controls symmetric to u obtained by applying one or more of three transformations defined below. These transforms are isometries; if u is optimal from q , then the transformed control $\tau(u)$ is optimal from the transformed configuration $T(q)$, where τ is some combination of τ_1 , τ_2 , and τ_3 , and T is the corresponding combination of T_1 , T_2 , and T_3 .

Geometrically, the transformations reflect the plane across the origin or across one of three other lines: the x -axis, a line Δ_θ at angle $(\pi + \theta_s)/2$, or the line Δ_θ^\perp at angle $\theta_s/2$. The three transformations on controls are:

τ_1 :	Swap \uparrow and \downarrow
τ_2 :	Reverse order
τ_3 :	Swap \curvearrowright and \curvearrowleft

The three corresponding transformations on trajectories are:

T_1 :	$q = (-x, -y, \theta)$
T_2 :	$(x, y) = \text{Rot}(\theta)(x, -y)$
T_3 :	$q = (x, -y, -\theta)$

Each transformation is its own inverse, and the three transformations commute. For any given base trajectory, the transformations yield up to seven different symmetric trajectories. The result is that all optimal trajectories fall in one of nine symmetry classes. We sort the optimal trajectory structures presented in table 1 by symmetry class to obtain table 2. We can analyze all types of trajectories by analyzing just the ‘base’ trajectories (the first column of table 2), and then applying the transformations T_1, T_2, T_3 to obtain the other members of the class.

10 Time optimal trajectories

In this section we identify the time optimal trajectories between any given start and goal configuration. We use a goal-centric coordinate system, with the origin coincident with the goal position, and the x axis aligned with the goal heading.

The symmetries of the previous section greatly simplify our analysis. We analyze only a base subset of the possible starting configurations, and use the symmetries to extend the analysis to the other starting configurations. In principle, the analysis is carried out by the following steps:




Figure 10: Given an optimal trajectory from “base” with heading θ_s to the origin with heading $\theta_g = 0$, transformations T_1, T_2 , and T_3 yield up to seven other optimal trajectories symmetric to the original.

1. For each trajectory type, we identify every feasible choice of start configuration (x, y, θ) . This defines a map from trajectory type to a region of configuration space.
2. Now we consider a point in configuration space (x, y, θ) . If it is in only one region, then the corresponding trajectory type is optimal from that point.
3. When regions overlap, we derive additional necessary conditions for optimality or calculate the actual times for each trajectory type to disambiguate.

To illustrate this procedure, we present the following example (figure 11). The feasible regions for $\downarrow\curvearrowright\uparrow\curvearrowleft$ and $\curvearrowleft\downarrow\curvearrowright\uparrow$ overlap. For almost all q_s in the overlap, there are two possible extremals but only one true optimal trajectory. It turns out that the Δ_θ line is a decision boundary: for q_s to the right of Δ_θ the optimum is $\curvearrowleft\downarrow\curvearrowright\uparrow$, and to the left of Δ_θ the optimum is $\downarrow\curvearrowright\uparrow\curvearrowleft$. Figure 11 illustrates the proof. First we observe that the alternatives give equal time on the Δ_θ line, because that line is the axis of reflection for the $T_1 \circ T_2$ isometry. So *both* trajectories are optimal on Δ_θ .

Consider a $\downarrow\curvearrowright\uparrow\curvearrowleft$ trajectory from the start pose shown. When the trajectory crosses Δ_θ during the \downarrow action, the remaining cost is unchanged if it switches to $\curvearrowleft\downarrow\curvearrowright\uparrow$. But then the total trajectory would have a structure of $\downarrow\curvearrowleft\downarrow\curvearrowright\uparrow$, and would not be a legitimate extremal. \square

Similar techniques can be applied to the other regions. The end result is a mapping that defines for each point in configuration space the set of optimal trajectories from that point to the origin. This mapping is illustrated by showing a slice at $\theta = \pi/4$ (figure 12). The mapping from start configuration to optimal trajectory is usually, but not always, unique. At some boundaries in the figures there are two distinct trajectories that give the same time cost. More interesting is the case at $\theta = 0$ where a continuum of different trajectories of type A are all optimal, bounded by optimal trajectories of type B.




Figure 11: An example of overlapping regions. (a) Plots of the feasible regions for $\curvearrowright\downarrow\curvearrowright\uparrow$ and $\downarrow\curvearrowright\uparrow\curvearrowright$. (b) The trajectory shown is extremal, but not optimal, so Δ_θ is the decision boundary. (c) The overlap region is partitioned according to which trajectory class is time optimal.

11 Algorithm for optimal control and value function




Figure 12: Optimal control for start configuration $q_s = (x, y, \frac{\pi}{4})$ and goal configuration $q_g = (0, 0, 0)$, with isocost lines. Coordinates are measured in units of b .

We now present an algorithm to determine the optimal trajectories between a given start and goal position, and the time cost of those trajectories. For each optimal trajectory structure, the necessary conditions yield a region. (Twelve such regions are shown in figure 12.) The algorithm determines which region(s) the start configuration (x, y, θ) falls in, and then calculates the value function for one of the optimal trajectory structures. For example, the function ValueBaseTSTS below calculates the cost of the fastest trajectory with a structure of $\curvearrowright\downarrow\curvearrowright\uparrow$.

```

Procedure ValueBaseTSTS( $q = (x, y, \theta)$ )
   $\arccos(1 - y) - \theta/2 - x + \sqrt{y(2 - y)}$ 
End ValueBaseTSTS

```

```

Procedure ValueBaseSTS( $q = (x, y, \theta)$ )
  If  $y = 0$  then  $|x| + \theta/2$ 
  else  $y(1 + \cos(\theta))/\sin(\theta) - x + \theta/2$ 
End ValueBaseSTS

```

```

Procedure ValueBaseTST( $q = (x, y, \theta)$ )
   $r = \|(x, y)\|$ 
   $\zeta = \arctan(y, x)$ 
   $r + \min(|\zeta| + |\zeta - \theta|, 2\pi + |\zeta| - |\zeta + \theta|)$ 
End ValueBaseTST

```

We now can define OptBVDD (optimal bounded velocity differential drive). The function recursively applies symmetry transforms until the configuration is in the base region.

The optimal trajectory structure can then be determined based on the necessary conditions for extremal trajectories to be optimal. The value for that trajectory structure is calculated. The recursion applies the appropriate combination of τ_1 , τ_2 , and τ_3 transforms to the base trajectory structure to determine the actual optimal trajectory structure.

```

Procedure OptBVDD( $q = (x, y, \theta)$ )
  if  $\theta \in (\pi, 2\pi)$  then  $\tau_3(\text{OptBVDD}(T_3(q)))$ 

   $r = \|(x, y)\|$ 
   $\zeta = \arctan(y, x)$ 
  if  $\zeta \in ((\theta + \pi)/2, \pi) \cup ((\theta - \pi)/2, 0)$ 
    then  $\tau_2(\text{OptBVDD}(T_2(q)))$ 
  if  $y < 0$  then  $\tau_1(\text{OptBVDD}(T_1(q)))$ 

  if  $\zeta \leq \theta$ 
    return( $\curvearrowright\downarrow\curvearrowleft$ , ValueBaseTST( $q$ ))
  else if  $y \leq 1 - \cos(\theta)$ 
    return( $\downarrow\curvearrowright\uparrow$ , ValueBaseSTS( $q$ ))
  else if  $r \geq \tan(\zeta/2)$ 
    return( $\curvearrowright\downarrow\curvearrowleft$ , ValueBaseTST( $q$ ))
  else
    return( $\curvearrowright\downarrow\curvearrowleft\uparrow$ , ValueBaseTSTS( $q$ ))
End OptBVDD

```

For the sake of brevity, certain special cases have

been omitted from the pseudocode presented. Whenever two symmetric regions share a boundary, the fastest trajectories for both regions are optimal on the boundary. For example, if the robot starts at $(0, 1, \pi)$, then both the trajectories $\curvearrowright\downarrow\curvearrowleft$ and $\curvearrowleft\uparrow\curvearrowright$ are optimal.

There are two other cases where multiple trajectories will be optimal. When $\theta_s = 0$, there may be a continuum of optimal five-action trajectories, bounded by two different four-action trajectories. When $\theta_s = \pi$, there will be a continuum of optimal $\tau_{\pi s}$ trajectories (Class E), bounded by two-action trajectories of class C.


In all cases, the above algorithm will return a single optimal trajectory. Some additional bookkeeping would allow all of the optimal trajectories to be returned.

Slices of this value function allow the regions in which various extremal trajectories are optimal to be seen more clearly. For example, figure 12 shows a slice for which the angle between the start and goal robot is fixed at $\pi/4$.


12 Performance limits

The level sets of the value function show the reachable configurations of the robot for some given amount of time. Figure 13a shows the shape of this region for time 1. (x , y , and time are normalized by b , the width of the robot.) Notice that in time 1, there are headings which the robot cannot reach; there is not enough time for the robot to turn to the heading $\theta = \pi$. The thin sliver-like appearance of the reachable configurations confirms the intuitive observation that it is easier to move forwards or backwards than sideways.


Figures 13b and 13c show the shapes of this region for time 2 and time 3 respectively. In the figure for time 2, there is enough time for the robot to reach any heading. In the figure for time 3, we can see that the cross sections for fixed θ are fatter; *i.e.*, with more time the robot is able to go nearly as far in the sideways directions as in the forwards or backwards directions. Figure 13d shows a cross sectional view of the configurations the robot can reach in times




(a)



(b)



(c)



(d)

Figure 13: Reachable configurations in normalized time 1, 2, and 3 (a, b, and c respectively), and a cross section for values 1, 1.5, 2, 2.5, and 3 (d).

1, 1.5, 2, 2.5, and 3.

13 Conclusion

A simple kinematic model of a differential drive robot with bounds on the speeds of the wheels was derived. We applied Pontryagin's maximum principle to derive geometric constraints on the optimal trajectories; trajectories satisfying these constraints are *extremal* trajectories. We derived conditions necessary for extremal trajectories to be optimal. These conditions require that optimal trajectories fall in one of 40 extremal trajectory classes. We then applied symmetries developed by Souères and Boissonnat, reducing the set of trajectory classes to be analysed to nine. We analysed each of the nine classes to determine the start and goal configurations for which it was optimal. This yields a simple algorithm to determine the optimal trajectory structure and cost between any two configurations.

Acknowledgements.

We would like to thank Jean Paul Laumond for guidance. We would also like to thank Al Rizzi, Howie Choset, Ercan Acar, and the members of the Manipulation Lab for helpful comments. We thank the Department of Energy for a Computational Science Graduate Fellowship, and the National Science Foundation for its support through grants 9900322 and 0082339.