Nearly Time-Optimal Paths for a Ground Vehicle

David A. Anisi† Johan Hamberg‡ Xiaoming Hu†
†Optimization and Systems Theory
Royal Institute of Technology
SE-10044 Stockholm, Sweden.

‡Department of Autonomous Systems
Swedish Defence Research Agency
SE-172 90 Stockholm, Sweden.

Abstract

It is well known that the sufficient family of time-optimal paths for both Dubins' as well as Reeds-Shepp's car models consist of the concatenation of circular arcs with maximum curvature and straight line segments, all tangentially connected. These time-optimal solutions suffer from some drawbacks. Their discontinuous curvature profile, together with the wear and impairment on the control equipment that the bang-bang solutions induce, calls for "smoother" and more supple reference paths to follow. Avoiding the bang-bang solutions also raises the robustness with respect to any possible uncertainties.
In this paper, our main tool for generating these "nearly time-optimal", but nevertheless continuous-curvature paths, is to use the Pontryagin Maximum Principle (PMP) and make an appropriate and cunning choice of the Lagrangian function. Despite some rewarding simulation results, this concept turns out to be numerically divergent at some instances. Upon a more careful investigation, it can be concluded that the problem at hand is nearly singular. This is seen by applying the PMP to Dubins' car and studying the corresponding two point boundary value problem, which turn out to be singular. Realizing this, we are able to contradict the widespread belief that all the information about the motion of a mobile platform lies in the initial values of the auxiliary variables associated with the PMP.

1 Introduction

In his pioneering work from 1957 [3], L.E. Dubins considered a problem which later was interpreted as finding the shortest continuously differentiable path between two given points taken by a car, for which the starting and ending directions are specified. In addition, the car is required to move with unit speed and subject to a minimum turning radius constraint. The car model considered, is

the unicycle robot model

$$\begin{aligned} \dot{x} &= v\cos\varphi \\ \dot{y} &= v\sin\varphi \\ \dot{\varphi} &= \omega, \end{aligned} \tag{1}$$

where the reference point $(x, y) \in \mathbb{R}^2$ is taken as the midpoint of the car's rear axle, while φ is used to denote the car's orientation angle. Further we let v and ω denote the linear and lateral control inputs respectively. Model (1) suggests that the car model is subject to a nonholonomic constraint

$$\dot{x}\sin\varphi - \dot{y}\cos\varphi = 0.$$

In the same paper, Dubins showed that every time-optimal path interconnecting any two given points in the state space, is the concatenation of circular arcs with maximum curvature and straight line segments, all tangentially connected. This basic path segments' duration and mutual order however is a much more delicate matter.
In 1990, J.A. Reeds and L.A. Shepp extended these results to the case when the car is augmented with a reverse gear [5]. Although appearing insignificant, allowing the car to perform backwards motions as well, turns out to have implications on various issues, including the car's controllability and symmetry properties [7]. The time-optimal paths will however still be of the same type as those accosiated with Dubins' car, i.e. essentially bang-bang solutions.

The abovementioned bang-bang solutions, suffer from some drawbacks. Their discontinuous curvature profile, together with the wear and impairment on the control equipment that the bang-bang solutions induce, calls for "smoother" and more supple reference paths to follow. Stated more precisely, we would like the generated paths to have continuous and bounded curvature (i.e. $\in C^2$). Important contributions to the study of this problem have been given by Boissonnat *et. al* in [2], where they proved that all candidates for optimality are made of the concatenation of circular arcs, straight line segments and clothoid pieces. Unfortunately they also showed that the optimal paths consist of infinity many pieces whenever they contain a straight line segment, which clearly occurs whenever the configurations to be connected are far apart. Avoiding this, Sheuer and Laugier have considered the solution when the number of switchings are restricted to a finite number [6].

Our approach is to investigate how tools from Optimal Control and above all the Pontryagin Maximum Principle (PMP) can be used for generating such nearly time-optimal paths. We thus consider the following Optimal Control

Problem:

minimize $\int_0^T \mathcal{L} \mathrm{d}t$	
subject to	
$\dot{x}(t) = [v(t)\cos\varphi(t), v(t)\sin\varphi(t), \omega(t)]^T$	system dynamics
$x(0) = [x_i, y_i, \varphi_i]^T = X_i$	initial configuration
$x(T) = [x_f, y_f, \varphi_f]^T = X_f$	final configuration
$x(\cdot) \in \mathcal{X} = \mathbb{R}^2 \times S^1$	state constraint
$v(\cdot) \in [-1, 1]$	linear velocity constraint
$\omega(\cdot) \in [-1, 1]$	lateral velocity constraint
$\mathcal{L}(x, v, \omega) : \mathcal{X} \times \mathbb{R}^2 \to \mathbb{R}$	integral cost function (Lagrangian)
$T \in (0, \infty)$	final time (to be minimized).

The cunning choice of an appropriate Lagrangian function, $\mathcal{L}$, should clearly be made with some practical, control-motivated reasons in mind. What objectives to consider when making that choice and examples of suitable Lagrangian functions are presented in section 2. A nearly time-optimal, but nevertheless smoother and more supple solution, also brings more flexibility and robustness, in the presence of uncertainty. All the above mentioned advantages are illustrated in section 3 where some simulation results are presented. It turns out however, that the presented concept suffers from numerical instability properties. The origin of this undesirable behavior is revealed in section 4.

2 Selection of the Lagrangian

In this section we discuss how to make a cunning choice for an arbitrary Lagrangian function, $\mathcal{L}$, with some practical, control-motivated reasons in mind. We have three main objectives that should be reflected in the Lagrangian:

1. We have the obvious intention of reaching the terminal configuration as fast as possible. Hence, we wish to eliminate the solution candidates, in which unnecessarily amount of time is being wasted. This target is reached by including an *integral constant*, $\mathcal{L}_1 = 1$ in the integral cost function, $\mathcal{L}$.
2. In order to avoid the drawbacks of time-optimal, bang-bang solutions, we wish to obtain smoother paths. This issue involves introducing a *penalty* for steering the car with control values close to their boundary points. The penalty function enables us to encourage a more moderate driving style. The most intuitive contribution of this objective to $\mathcal{L}$, is any even polynomial, for instance $\mathcal{L}_2 = u^2$, where u refers to the linear and the lateral control indistinctly, so that $u \in \{v, \omega\}$.
3. Technological limitations, such as the car's turning radius, normally leave us with a restricted control domain. Our third and last objective when selecting $\mathcal{L}$, is to handle input saturations in a convenient manner. A direct approach, involves adopting an integral cost function, having the form of a well. For instance

$$\mathcal{L}_3 = \begin{cases} 0 & \text{if } u \in [-1,\ 1] \\ \infty & \text{else} \end{cases}$$

provides a solid *barrier*, in which the control function is allowed to take its values. This idea would be exploited in our starting choice of $\mathcal{L}$. However, as illustrated in the up-following discussion, by making a more clever selection of $\mathcal{L}$, we show a more subtle and implicit way to handle input restrictions.

The naive and adverse way to bond the contribution of our three objectives together, is to simply add them up, that is to set $\mathcal{L} = \sum_{i=1}^{3} \mathcal{L}_i$. We might however observe that by using superposition, the penalizing properties of $\mathcal{L}_2$ can be welded together with the off-barriering characteristics of $\mathcal{L}_3$ and be expressed by means of a common cost function, $\mathcal{L}_0$.
In order to obtain even more versatility, let $\mathcal{L}_0$ be proportional to what we call a *design parameter*, ε. Then by adjusting ε, we will be able to decide how close to the time-optimal solution we wish to find ourselves. Setting $\varepsilon = 0$ yields $\mathcal{L} = \mathcal{L}_1 = 1$ and thus equals this with the time-optimal problems considered in [3] and [5]. Then by gradually tuning ε up, we penalize the usage of the boundary points of the control domain increasingly, and hence put more value into a moderate driving style. This of course occurs at the expense of time-optimality.

We study the construction of nearly time optimal paths by starting with the following intuitively justified candidate for an appropriate Lagrangian function,

$$\mathcal{L} = \mathcal{L}_0 + \mathcal{L}_1 = \frac{\varepsilon}{2\left(1 - u^2\right)} + 1 \tag{2}$$

which is sketched in figure 1 and undeniably possesses both the desired penalizing and restricting properties. We will see, that this choice of integral cost function, will lead to analytically solvable equations for the optimal control. However, by means of a simple approximation trick, we will proceed and get extremely perspicuous and tidy expressions for the optimal control, simultaneously possessing the desired qualities we seek. This approximation then, will serve as a source of inspiration for the discussion to follow, where we derive the same results without employing the approximation step.

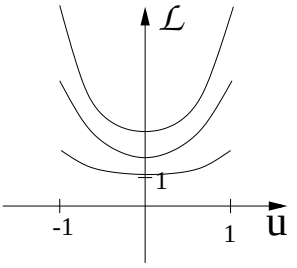

Figure 1: The integral cost function, sketched for different values on ε.

The Hamiltonian in this case becomes

$$\mathscr{H} = \lambda^T f - \mathcal{L} = \lambda_x v \cos\varphi + \lambda_y v \sin\varphi + \lambda_\varphi\, \omega - \frac{\varepsilon}{2\left(1 - v^2\right)} - \frac{\varepsilon}{2\left(1 - \omega^2\right)} - 1$$

Maximizing this Hamiltonian pointwise, involves finding the roots of its partial derivatives

$$\frac{\partial \mathscr{H}}{\partial v} = \sigma - \frac{\varepsilon v}{(1-v^2)^2} = 0 \quad (3)$$

$$\frac{\partial \mathscr{H}}{\partial \omega} = \lambda_\varphi - \frac{\varepsilon \omega}{(1-\omega^2)^2} = 0, \quad (4)$$

where $\sigma = \lambda_x \cos\varphi + \lambda_y \sin\varphi$. By virtue of the similarities between equations (3) and (4), it suffices to study either one. Consider then the partial derivate w.r.t. the lateral velocity and initially assume $\lambda_\varphi \neq 0$ (the $\lambda_\varphi \equiv 0$ case, will be examined later). Then, we are able to write equation (4) as

$$\Delta(\omega) := (1-\omega^2)^2 - \frac{\varepsilon}{\lambda_\varphi}\omega = 0. \quad (5)$$

Apparently (5) would only have real roots on the $\text{sign}(\lambda_\varphi)\omega \geq 0$ half axis. Since $\Delta(0) > 0$, $\Delta(\text{sign}(\lambda_\varphi)1) < 0$, and $\text{sign}(\lambda_\varphi)\Delta'(\omega) < 0$ for all $\text{sign}(\lambda_\varphi)\omega \in (0,\ 1)$, we conclude that only one root lies in the admissible interval, or equivalently that the pointwise maximization of the Hamiltonian has a unique solution. In figure 2 the real roots of equation (5) are plotted as the intersection points of the functions $f_1(\omega) = (1-\omega^2)^2$ and $f_2(\omega) = \frac{\varepsilon}{\lambda_\varphi}\omega$.

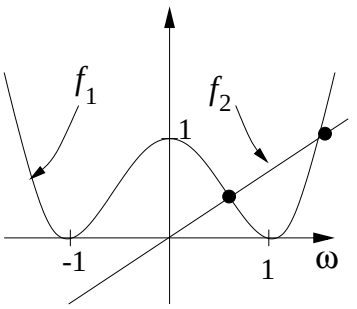

Figure 2: The roots of equation (5) correspond to the intersection points of $f_1(\omega) = (1-\omega^2)^2$ and $f_2(\omega) = \frac{\varepsilon}{\lambda_\varphi}\omega$.

In the time-optimal case ($\varepsilon = 0$), the slope of $f_2(\omega)$ is zero, hence the two possible values for ω^* are ± 1. This naturally corresponds to the minimum radius left- and right turn respectively. This holds for all non-zero values on λ_φ. But as we set $\varepsilon \neq 0$, only the quadrant of location of our only real and admissible root, is determined by $\text{sign}\lambda_\varphi$. Its exact value is indefinite but nevertheless, continuously varying with λ_φ. We further note that the boundary values of the lateral control, are not reached unless $\lambda_\varphi \to \infty$, excluding the occurrence of bang-bang solutions. These are very desirable properties.

Now, what happens as $\lambda_\varphi \to 0$ and we cross the border line? Will ω^* remain continuous even at this point? Since $\lambda_\varphi = 0$, $\omega = 0$ is a solution to equation (4), and we have shown that only one root lies in the admissible interval when $\lambda_\varphi \neq 0$, one can easily use the implicit function theorem to show that ω^* varies continuously as λ_φ evolves and changes sign.

Although it is fully possible to analytically solve equation (4) for ω^*, the resulting expressions looks anything but tidy. Therefore, we approximate $f_1(\omega)$, $\omega \in [-1, 1]$ in figure 2, with the upper half of a unit circle. With this approach, expressing the lateral control as a function of ε, reduces to finding the point of

intersect between a straight line and a semi-circle. Referring to figure 3 and considering the fact that $\tan\vartheta$ equals the direction of the straight line $\frac{\varepsilon}{\lambda_\varphi}$, we get

$$\omega^* = \cos\vartheta = \frac{1}{\sqrt{\tan^2\vartheta + 1}} = \frac{1}{\sqrt{(\frac{\varepsilon}{\lambda_\varphi})^2 + 1}} = \frac{\lambda_\varphi}{\sqrt{\varepsilon^2 + \lambda_\varphi^2}}. \tag{6}$$

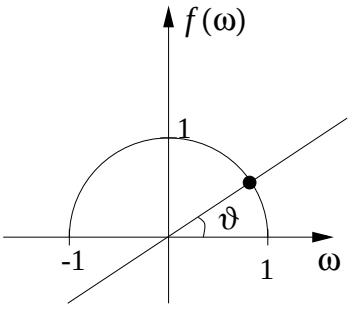

Figure 3: The circle segment approximates $f_1(\omega)$, $\omega \in [-1, 1]$ in figure 2 and leads to some rewarding results.

We observe that by setting $\varepsilon = 0$ in equation (6), we get $\omega^* = \text{sign}\,\lambda_\varphi$, which is to be recognized as the optimal control associated with the time-optimal case. Then by gradually tuning the design parameter up, we damp the fluctuating behavior of the optimal control whenever λ_φ changes sign, and thereby replace the bang-bang solutions with smoother and more pliable paths (cf. figure 4).

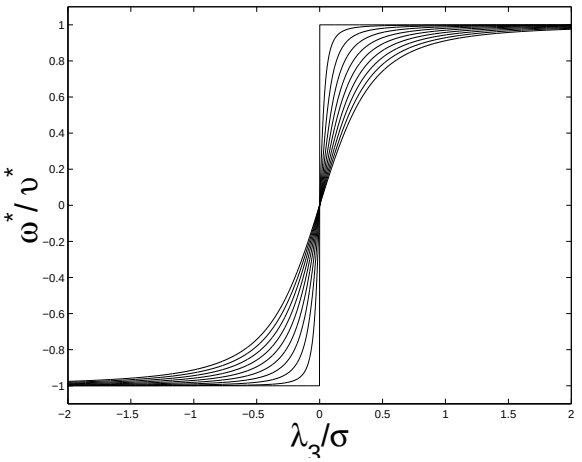

Figure 4: The ε-dependence of the optimal control. ε varying between 0 and 0.5.

Analogous reasoning applies for the linear control v^*, with σ playing the role of λ_φ. The summarizing results of this example will therefore be

$$\boxed{\begin{array}{lcl} v^* & = & \frac{\sigma}{\sqrt{\varepsilon^2+\sigma^2}} \\ \omega^* & = & \frac{\lambda_\varphi}{\sqrt{\varepsilon^2+\lambda_\varphi^2}} \end{array}}$$

Nearly time-optimal paths

Motivated by the perspicuous results and their desirable properties obtained in the solution to the optimal control problem defined by (2), we make the following choice for Lagrangian

$$\mathcal{L} = 1 - \varepsilon\sqrt{1-u^2}, \tag{7}$$

its circular form, evidently influenced by the approximation made in the previous example. The Hamiltonian then becomes

$$\mathscr{H} = \lambda^T f - \mathcal{L} = \lambda_x v\cos\varphi + \lambda_y v\sin\varphi + \lambda_\varphi\,\omega + \varepsilon\sqrt{1-v^2} + \varepsilon\sqrt{1-\omega^2} - 1 \tag{8}$$

and its pointwise maximization gives

$$\frac{\partial\mathscr{H}}{\partial v} = \sigma - \frac{\varepsilon v}{\sqrt{1-v^2}} = 0 \implies \sqrt{1-v^2} = \frac{\varepsilon}{\sigma}\,v \tag{9}$$

$$\frac{\partial\mathscr{H}}{\partial \omega} = \lambda_\varphi - \frac{\varepsilon\omega}{\sqrt{1-\omega^2}} = 0 \implies \sqrt{1-\omega^2} = \frac{\varepsilon}{\lambda_\varphi}\,\omega, \tag{10}$$

where $\sigma = \lambda_x\cos\varphi + \lambda_y\sin\varphi$. The solutions of both (9) and (10), involve finding the point of intersection between a straight line, with a known slope ($\frac{\varepsilon}{\sigma}$ or $\frac{\varepsilon}{\lambda_\varphi}$) and the upper half of a unit circle and is best interpreted by means of figure 3. In accordance with the calculations made previously, we obtain

$$\boxed{\begin{array}{rcl} v^* &=& \frac{\sigma}{\sqrt{\varepsilon^2+\sigma^2}} \\ \omega^* &=& \frac{\lambda_\varphi}{\sqrt{\varepsilon^2+\lambda_\varphi^2}} \end{array}} \tag{11}$$

Reflecting upon what this choice of the Lagrangian has resulted in, we are able to pinpoint some distinguished characteristics; firstly, notice that $\mathcal{L}_0$ in this case, conceptually differs from the ones considered so far. This one merely satisfies our penalizing objective, while no solid control restricting barrier, is imposed. Nevertheless, we can conclude from the expressions for the optimal controls (equation (11)), that even this requirement is met. For all non-zero values on the design parameter ε, $|u^*| < 1$, which is the saturation boundary for the optimal control. This provides us with a more subtle and implicit approach for handling control constraints.
Secondly, we observe that we have the possibility to make a continuous and arbitrary adjustment of the design parameter, in order to damp the fluctuating behavior of the optimal control pertaining to time-optimal solutions. This introduced flexibility, is illustrated in figure 4 and enables us to designedly avoid time-optimal solutions and their drawbacks while obtaining more supple paths. Yet another advantage is the smoothing effect of u^* on the systems Hamiltonian function H. This benefit is to be discussed more thoroughly in section 4.

3 Simulations and experiments

In addition to the systems dynamics that governs the time evolution of the state vector and the optimal control law specified by equation (11), the time evolution of the auxiliary variables will be of interest when we carry out our simulations. The TPBVP that we ought to consider is

Two Point Boundary Value Problem (TPBVP)

$$\begin{cases} \dot{x} = \frac{\partial H}{\partial \lambda_x} = v^* \cos\varphi \\ \dot{y} = \frac{\partial H}{\partial \lambda_y} = v^* \sin\varphi \\ \dot{\varphi} = \frac{\partial H}{\partial \lambda_\varphi} = \omega^* \end{cases} \qquad \begin{cases} \dot{\lambda_x} = -\frac{\partial H}{\partial x} = 0 \\ \dot{\lambda_y} = -\frac{\partial H}{\partial y} = 0 \\ \dot{\lambda_\varphi} = -\frac{\partial H}{\partial \varphi} = v^*[\lambda_y \cos\varphi - \lambda_x \sin\varphi] \end{cases}$$

$$\text{s.t. } x(0) = X_i \text{ and } x(T) = X_f \tag{12}$$

where

$$\boxed{\begin{array}{lcl} v^* & = & \frac{\sigma}{\sqrt{\varepsilon^2+\sigma^2}} \\ \omega^* & = & \frac{\lambda_\varphi}{\sqrt{\varepsilon^2+\lambda_\varphi^2}} \end{array}}$$

Since the state variables are specified at both the initial and final time instant, while the auxiliary variables are unrestricted, we have a mixed boundary value problem at hand. Taking resource in the numerical methods, the so called *shooting method* is the most well-tried and dependable technique.

Figure 5 is the outcome of our first simulation and serves as a perfect example of how choosing $\mathcal{L}$ as prescribed by equation (7) results in such expressions for optimal input that meets all our three objectives discussed previously. The task is to steer the car between $X_i = [-2,\ 2,\ -\frac{\pi}{2}]$ and $X_f = [0,\ 0,\ -\frac{\pi}{2}]$. The time-optimal path is obtained by setting $\varepsilon = 0$ and is sketched with a thicker line in figure 5. However, by tuning ε up gradually, we are able to digress from this time-optimal (bang-bang) solution in a continuous and controlled manner.

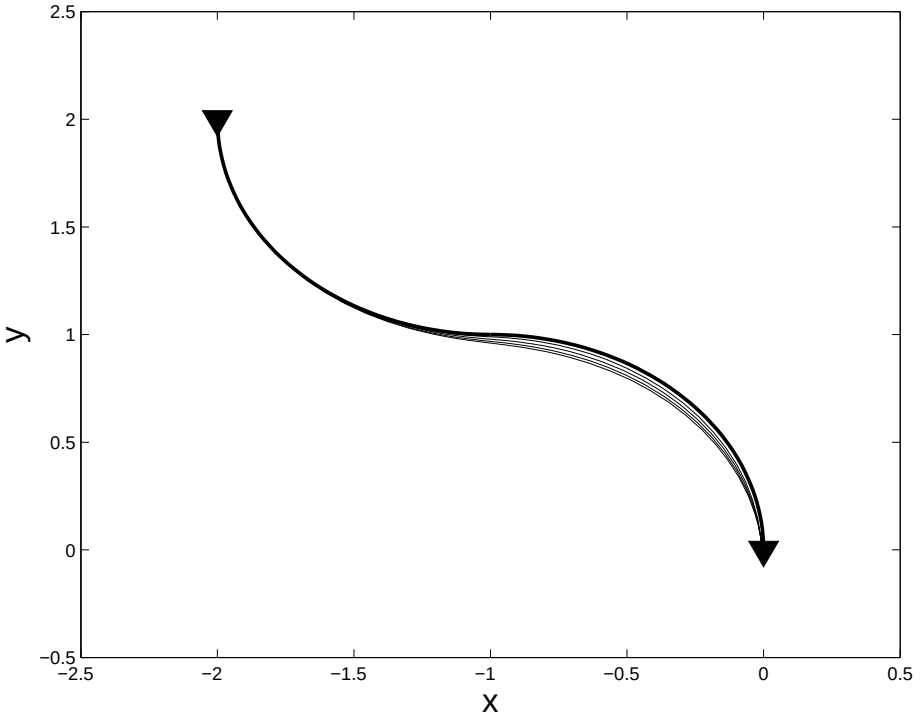

Figure 5: By tuning ε up gradually, we digress from the time-optimal (bang-bang) solution.

The embedded robustness in these more supple paths is illustrated in our next simulation where we wish to study the flexibility of the mobile platform with

respect to a change in a *prescribed* time of arrival, $\widehat{T}$. Such flexibility is of great importance when simultaneous rendezvous problems for a team of platforms are considered.

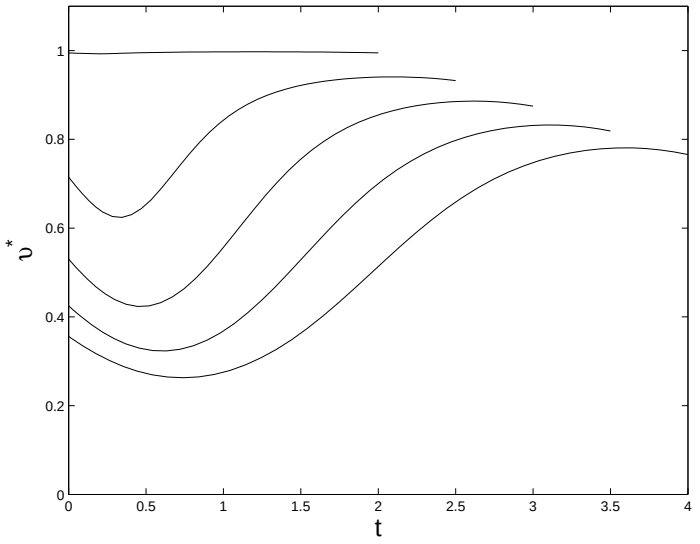

Figure 6: As the time of arrival sets to higher values, a more moderate driving style is being adopted. By decreasing the value of the linear velocity, the car "wastes time" and is thereby able to adjust its time of arrival within a considerable interval of time.

Referring to figure 6 and 7, the task in all five trials is to steer between the same two prescribed configurations. However, the prescribed time of arrival, $\widehat{T}$, varies. As seen in figure 6, when $\widehat{T}$ is set to 2, the linear velocity $v^*(t)$, almost takes its highest value, i.e. equals 1, during the entire time interval - this in order to be able to make it to the final configuration at the prescribed time of arrival. But as we set $\widehat{T}$ to higher and higher values, we note how the linear velocity decreases and a more moderate driving style is being adopted.
The corresponding control in the the lateral direction $\omega^*(t)$, can be seen in figure 7, where we note that for $\widehat{T} = 2$, $\omega^*(t)$ is relatively bang-bang. But as the value of the arrival time increases, the fluctuating behavior of the angular velocity is being damp, resulting in much smoother paths.

These two abovementioned simulations, illustrate that the proposed control law meets all our requirements and fulfills all our objectives thus far. However, the numerical computations in the shooting method, turns out to be divergent at some instances. In the next section, we are to locate the origin of this undesirable behavior.

4 The Singular Property of the Problem

In order to get a glimpse beneath the surface and gain some insight about the reasons for the shooting method to diverge, we recur to Dubins' problem. The Hamiltonian function in this case becomes (cf. [1])

$$H(x,\lambda) = \mathscr{H}(x,\lambda,\omega^*) = \lambda_x \cos\varphi + \lambda_y \sin\varphi + |\lambda_\varphi| - 1. \tag{13}$$

Since neither of the state variables specifying the position of the platform (i.e. x or y) is included in H, the corresponding auxiliary variables, λ_x and λ_y, are

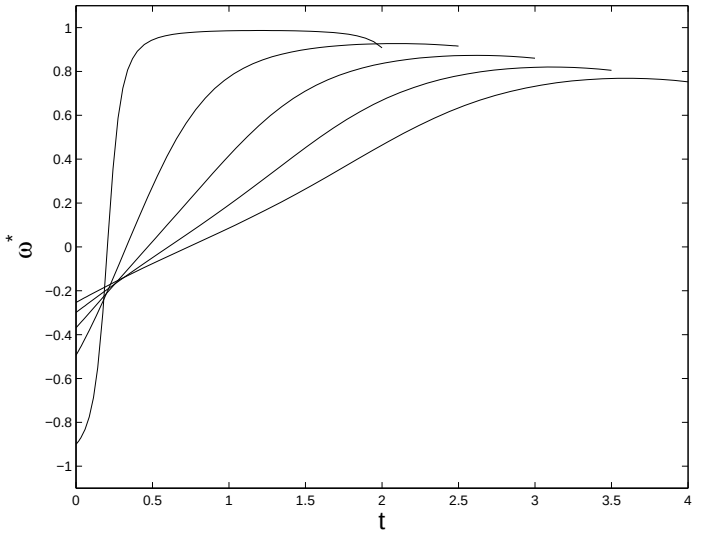

Figure 7: As the time of arrival is postponed, the generated paths become smoother.

cyclic, that is they are time constants. Then by setting

$$\begin{cases} \lambda_x = -\mu \cos\varphi_0 \\ \lambda_y = \quad \mu \sin\varphi_0, \end{cases}$$

we are able to write equation (13) as

$$H = \mu \cos(\varphi - \varphi_0) + |\lambda_\varphi| - 1.$$

Now, from the standard results regarding the PMP, it follows that the Hamiltonian is constant on a full trajectory - that is, the level curves of H correspond to different trajectories for the mobile platform. Let us therefore pay attention to these. The level curves of the system are sketched by means of *Mathematica* in figure 8. The reason for rewriting the cyclic auxiliary variables (λ_x and λ_y) in terms of μ and φ_0 might be more obvious now, since the constant φ_0 solely adjusts the horizontal alignment of the level curves, while the constant μ specifies their depth. We are thus able to study the full motion of the mobile platform by just paying attention to the time evolution of two variables, namely $\varphi(t)$ and $\lambda_\varphi(t)$. Generating a typical path for Dubins' car which consists of a circular movement, followed by a straight line movement followed by yet another movement on a circular arc is a singular problem and can not be achieved by just finding an appropriate value on $\lambda_\varphi(0)$. Notice that $\lambda_v(0)$ and $\lambda_\omega(0)$ do not effect the shape of the level curves.
Same conclusion can be drawn from studying figure 9 where the level surface of the Hamiltonian can be seen. We note that two smooth surfaces are seamed together at a joint, centered at the axis along $\lambda_\varphi = 0$. Hence the Hamiltonian has different derivatives on different sides of this joint, making the right hand side (the dynamics) of the TPBVP (12) a discontinuous function. Then it follows from standard results on differential equations, that in the case of a discontinuous dynamics, not even existence of a solution is guaranteed, even less its optimality or uniqueness. It is a widespread idea that all the information about the motion of a mobile platform lies in the initial values of the auxiliary variables $\lambda(0)$. We have however shown that this does not hold true in all cases and that a more

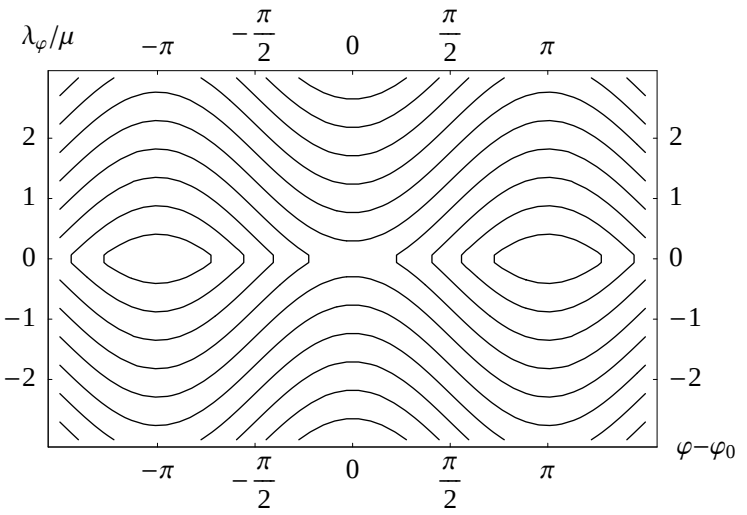

Figure 8: The level curves of the Hamiltonian correspond to different trajectories of the mobile platform. Generating a path for Dubins' car by choosing an appropriate $\lambda(0)$ is a singular problem.

careful analysis of the system properties must be carried out in order to be able to draw any conclusions about that matter.

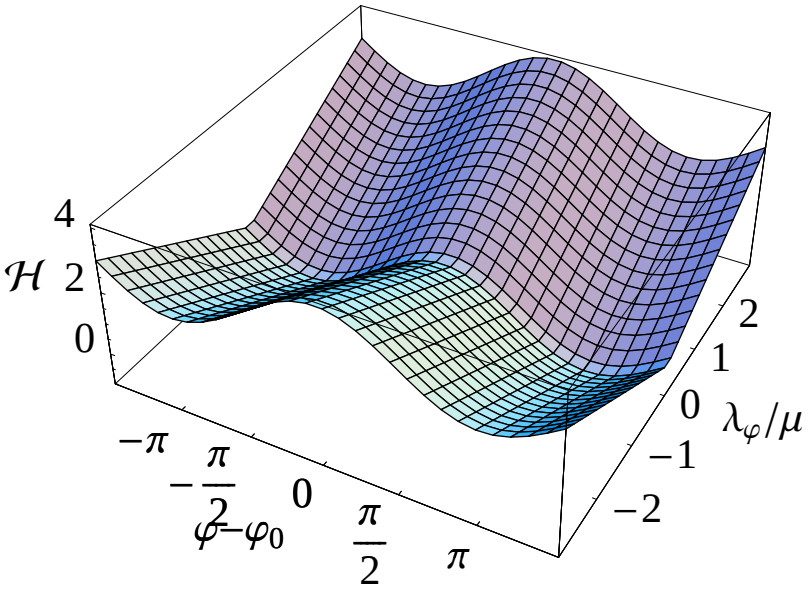

Figure 9: The level surface of the Hamiltonian function for Dubins' problem.

Let our next study involve the Hamiltonian when we utilize the control law specified by equation (11). Combining this with equation (8) yields the following expression for the Hamiltonian function

$$
\begin{aligned}
H_{\mathrm{CRS}}(x,\lambda) &= \mathscr{H}(x,\lambda,u^*) \\
&= \frac{\sigma^2}{\sqrt{\varepsilon^2+\sigma^2}} + \frac{\lambda_\varphi^2}{\sqrt{\varepsilon^2+\lambda_\varphi^2}} + \varepsilon\sqrt{1-\frac{\sigma^2}{\varepsilon^2+\sigma^2}} + \varepsilon\sqrt{1-\frac{\lambda_\varphi^2}{\varepsilon^2+\lambda_\varphi^2}} - 1 \\
&= \sqrt{\varepsilon^2+\sigma^2} + \sqrt{\varepsilon^2+\lambda_\varphi^2} - 1. \qquad (14)
\end{aligned}
$$

To fully understand the advantage with (14), let us present the expression for the Hamiltonian for the Reeds-Shepp car and make a comparison (cf. [1]).

$$H_{\mathrm{RS}}(x,\lambda) = \mathscr{H}(x,\lambda,u^*) = |\sigma| + |\lambda_\varphi| - 1$$

Comparing this with equation (14), we see that by introducing ε, we have abolished the discontinuous properties of H_{RS}. H_{CRS} do not have a sharp joint and is a continuously differentiable function. Possessing a continuously differentiable Hamiltonian is noticeable, since it improves and rectifies the numerical issues of the optimal control problem. However, our ambition is to produce *nearly time-optimal* paths, hence we are concerned with rather small values on the design parameter ε. As a result, the corresponding improvements in the numerical properties of the problem at hand, will be comparatively small. In order to get a significant improvement and obtain numerically stable algorithms for real-life applications, we have to tune ε up to considerable values. That in turn, penalizes the input signals accordingly and results in paths that differ too much from the time-optimal paths to be classified as "nearly time-optimal". Similar paths might more conveniently be generated by any "non-optimal" planning scheme, such as the cell-decomposition or visibility graph method (cf. [4]).

We have also investigated the possible approach of choosing $\mathcal{L}$ differently, so that better numerical properties are obtained for relatively low values on ε. A representative candidate for such a choice is $\mathcal{L} = \varepsilon\ln(1-u^2)+1$. The common characteristic of suitable Lagrangians for this approach, is that they put a considerable penalty on the input signal even for relatively low values on ε. The net outcome of the simulations adopting such Lagrangian is however similar to the ones aforementioned. There is obviously a trade-off between the numerical stability of the problem and the magnitude of the penalty imposed on the control function on one hand, and yet another one between the imposed penalty and the appearance of the generated paths. The bottom line is that, even in the case when we do possess a continuously differentiable H, it is not trivial to find appropriate starting values for $\lambda(0)$ which take us to the prescribed final configuration.

5 Summary and discussion

In this paper, we investigate how tools from Optimal Control and above all the Pontryagin Maximum Principle (PMP) can be used to generate nearly time-optimal, but nevertheless continuous-curvature paths for a ground vehicle. The unicycle robot model have been used to describe the kinematics of the vehicle. Initially we discussed what objectives to consider when making the choice for an arbitrary Lagrangian function and exemplified some appropriate choices that satisfies all our objectives. These qualities were then confirmed by the presented

simulation results. However, the numerical computations in the shooting method turn out to be divergent at some instances. Upon a more careful investigation, we could conclude that the problem at hand is nearly singular.
It is possible to rectify the numerical difficulties to some extend by either significantly increasing the value of the design parameter ε, or choosing another type of Lagrangian. The common characteristic of suitable Lagrangians for this line of action, is that they put a considerable penalty on the input signal even for relatively low values on ε. There is obviously a trade-off between the numerical stability of the problem and the magnitude of the penalty imposed on the control function on one hand, and yet another one between the imposed penalty and the appearance of the generated paths. Hence, the generated paths when striving to get a numerically stable algorithm for real-life applications, are too unlike the time-optimal bang-bang solutions, to be classified as "nearly time-optimal" paths and thus have to be disregarded. Similar paths might more conveniently be generated by any "non-optimal" planning scheme, such as the cell-decomposition or visibility graph method (cf. [4]).

When striving to localize the origin of the above mentioned undesirable behavior, we concluded that adopting shooting method on Dubins' problem is a singular problem. It is a widespread belief that all the information about the motion of a mobile platform lies in the initial values of the auxiliary variables associated with the PMP. We have however shown that this does not hold true in all cases and that a more careful analysis of the system properties must be carried out in order to be able to draw any conclusions about that matter.

References

[1] David A. Anisi. "Optimal Motion Control of a Ground Vehicle". Master's thesis, Royal Institute of Technology (KTH), Stockholm, Sweden, 2003.

[2] J.D. Boissonnat, A. Cerezo, and J. Leblond. "A note on shortest paths in the plane subject to a constraint on the derivative of the curvature". Research report 2160, Inst. Nat. de Recherche en Informatique et en Automatique, Jan. 1994.

[3] 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.

[4] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, 1991.

[5] 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.

[6] A. Scheuer and Ch. Laugier. "Planning Sub-Optimal and Continouous-Curvature Paths for Car-Like Robots". In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 25–31, 1998.

[7] Héctor J. Sussman and Guoqing Tang. "Shortest paths for the Reeds-Shepp car: a worked out example of the use of geometric techniques in nonlinear optimal control". Technical Report SYCON-91-10, Department of Mathematics, Rutgers University, 1991.