Discontinuous Feedbacks, Discontinuous Optimal Controls, and Continuous-Time Model Predictive Control

Fernando A. C. C. Fontes
Officina Mathematica
Departamento de Matemática para a Ciência e Tecnologia
Universidade do Minho,
4800-058 Guimarães, Portugal.
tel: +351 253 510 425, fax: +351 253 510 401
e-mail: ffontes@mct.uminho.pt, http://www.mct.uminho.pt/ffontes/

November 2002; to appear in IJRNC

Abstract

It is known that there is a class of nonlinear systems that cannot be stabilized by a continuous time-invariant feedback. This class includes systems with interest in practice, such as nonholonomic systems, frequently appearing in robotics and other areas. Yet, most continuous-time Model Predictive Control (MPC) frameworks had to assume continuity of the resulting feedback law, being unable to address an important class of nonlinear systems. It is also known that the open-loop optimal control problems that are solved in MPC algorithms may not have, in general, a continuous solution. Again, most continuous-time MPC frameworks had to artificially assume continuity of the optimal controls or, alternatively, impose some demanding assumptions on the data of the optimal control problem to achieve the desired continuity. In this work we analyse the reasons why traditional MPC approaches had to impose the continuity assumptions, the difficulties in relaxing these assumptions, and how the concept of "sampling feedbacks" combines naturally with MPC to overcome these difficulties. A continuous-time MPC framework using a strictly positive inter-sampling time is argued to be appropriate to use with discontinuous optimal controls and discontinuous feedbacks. The essencial features for the stability of such MPC framework are reviewed.

Key words. Predictive control; Receding horizon; Discontinuous feedbacks; Continuity of optimal controls, Nonlinear stability analysis, Nonholonomic systems.

1 Introduction

There are two continuity assumptions that are imposed in most continuous-time Model Predictive Control (MPC) approaches. Those are the continuity of the controls solving the open-loop optimal control problems, and the continuity of the constructed feedback law; see e.g. the survey [32] and references therein. These continuity assumptions, in addition to being very difficult to verify, are a major obstacle in enabling MPC to address a broader class of nonlinear systems. One reason is because some nonlinear systems cannot be stabilized by a continuous time-invariant feedback control [6, 41].

Discontinuous feedbacks are needed every time logical decisions have to be made (an example is to decide whether a robot should move to the left or to the right to avoid an obstacle it encounters in its path). Even motion in an unconstrained Euclidean space may give rise to such need. Systems that cannot be stabilized by a continuous feedback include some mechanical and nonholonomic systems with interest in practice, such as wheeled vehicles, robot manipulators, among many others ([26]).

Nonholonomic systems are, therefore, an important motivation (but not the only) to develop methodologies that allow the construction of discontinuous feedback controls. This justifies dedicating the next section describing this significant class of nonlinear systems. A section discussing discontinuous feedbacks follows.

A related issue, in MPC, is the study of discontinuous controls, i.e. discontinuous solutions to the optimal control problems being solved, and its use in the construction of the solution. This subject has some noteworthy features of independent interest that are explored in Section 4. We discuss some problems that arise when continuity of the optimal controls is assumed, as is usually the case in the existing continuous-time MPC literature.

The difficulties of relaxing the two continuity assumptions within MPC are studied along this work. We show how to overcome such difficulties and to integrate discontinuous controls and discontinuous feedbacks within an appropriate continuous-time MPC framework. One of the main features of the framework is to employ a strictly positive inter-sampling time. The convenient use of such inter-sampling time, allows the use of both discontinuous feedbacks and discontinuous controls. In addition, a positive inter-sampling time is much more natural to use in MPC , even within a continuous-time context. Optimization problems, certainly take some time to be solved, and thus it is not possible to solve them at every instant of time.

Several continuous-time MPC approaches using a positive inter-sampling time have been proposed (e.g. [30, 34, 35, 7, 15, 8, 24, 28]). Despite that, exploiting the benefits of such feature to relax the continuity assumptions has essentially been neglected in MPC literature.

The last sections discuss the main features required to guarantee stability for the MPC frameworks allowing discontinuities.

2 Nonholonomic Systems

Nonholonomic systems are, typically, completely controllable but instantaneously they cannot move in certain directions. Although these systems are allowed to reach any point in the state space and to move, eventually, in any direction, at a certain time or state there are constraints imposed on the motion – the so-called nonholonomic constraints. Some of the interesting examples are the wheeled vehicles, which, at a certain instant can only move in a direction perpendicular to the axle connecting the wheels.

Consider the mobile robot of a unicycle type in Fig. 1 which is represented by the model

$$\begin{cases} \dot{x} = (u_1 + u_2) \cdot \cos\theta \\ \dot{y} = (u_1 + u_2) \cdot \sin\theta \\ \dot{\theta} = u_1 - u_2 \end{cases}$$

where $u_1, u_2(t) \in [-u_{max}, u_{max}]$.

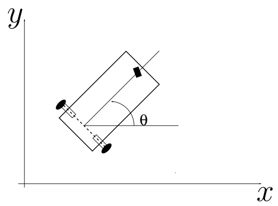

Figure 1: A unicycle mobile robot.

The coordinates (x, y) are the position in the plane of the midpoint of the axle connecting the rear wheels, and θ denotes the heading angle measured from the x-axis. The controls u_1 and u_2 are the angular velocity of the right and left wheels, respectively. If the same velocity is applied to both wheels, the robot moves along a straight line (maximum forward velocity when $u_1 = u_2 = u_{max}$). The robot can turn by choosing $u_1 \neq u_2$ (when $u_1 = -u_2 = u_{max}$ the robot turns anticlockwise around the midpoint of the axle).

The velocity vector is always orthogonal to the wheel axis: the nonholonomic constraint

$$(\dot{x}, \dot{y})^T (\sin\theta, -\cos\theta) = 0.$$

When trying to obtain a linearization of this system around any operating point $(x_0, y_0, \theta_0, u_1 = 0, u_2 = 0)$ we can easily conclude that the resulting linear system is not controllable. Therefore, linear control methods, or any auxiliary procedure based on linearization, cannot be used to handle this system.

There are several systems, many with very simple dynamics, possessing this nonholonomic characteristic. A rolling disc on a plane (a wheel) has a nonholonomic constraint preventing it from slipping sideways. Even a sphere allowed to roll freely on a plane has a nonholonomic constraint linking the angular motion with the linear motion.

A well-studied example is the nonholonomic integrator, also known as Brockett integrator, which is frequently mentioned to illustrate characteristics of the nonholonomic systems [6]:

$$\begin{cases} \dot{x} = u \\ \dot{y} = v \\ \dot{z} = yu - xv. \end{cases}$$

This system can be transformed into the unicycle system via changes of coordinates; see [39].

Another example we investigate here is the car-like vehicle (see Fig. 2). In this case, the heading is controlled by the angle of the two front directional wheels. In contrast with the previous example, this vehicle cannot turn with zero velocity; furthermore, it has a minimum turning radius. We represent this system by the model:

$$\begin{cases} \dot{x} = v \cdot \cos\theta \\ \dot{y} = v \cdot \sin\theta \\ \dot{\theta} = v \cdot c \end{cases}$$

with control inputs v and c satisfying $v \in [0, v_{max}]$ and $c \in [-c_{max}, c_{max}]$. The control v represents the linear velocity and c is the instant curvature (the modulus of c is the inverse of the turning radius; the minimum turning radius $R_{min} = |c_{max}|^{-1}$).

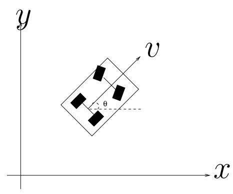

Figure 2: A car-like vehicle.

For this system, it is also easy to check that any linearization around the origin, or any other operating point, would lead to an uncontrollable system.

Another challenging characteristic of the nonholonomic systems is apparent in this example: it is not possible to stabilize it if just time-invariant continuous feedbacks are allowed. To illustrate this characteristic suppose the vehicle is on the positive side of the y-axis, heading towards the left-half plane, i.e. $x = 0$, $y > 0$, and $\theta = \pi$ (see Fig. 3). If $y \geq R_{min}$, then it is possible to move towards the origin following a semi-circle. So, the decision in this case would be to turn left. On the other hand, if the car is too close to the origin, $y < R_{min}$, then it needs to move away from the origin to obtain the correct heading. The decision in this case is to turn right. Therefore, the decision of where to turn, i.e. deciding input c, is a discontinuous function of the state — a discontinuous feedback.

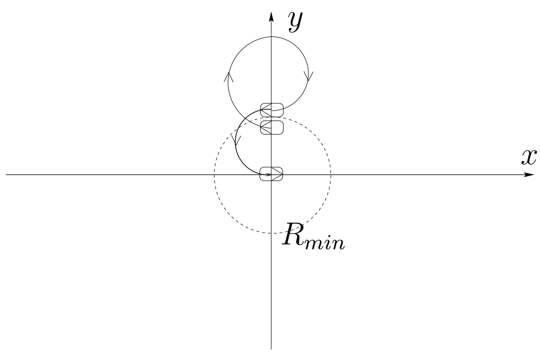

Figure 3: A discontinuous decision.

The need for discontinuous feedbacks to stabilize nonholonomic systems has been known for some time [6, 41]. The study of discontinuous feedbacks is the subject of the next section.

3 Discontinuous Feedbacks

As discussed in the previous section, allowing discontinuous feedbacks is essential to stabilize some nonlinear systems, in particular the nonholonomic systems. This need can be established rigourously, as follows.

Theorem 3.1 *(Brockett's Necessary Condition [6]) If the system $\dot{x} = f(x, u)$, with $u \in U$, admits a continuous stabilizing feedback, then for any neighbourhood $\mathcal{X}$ of zero, the set $f(\mathcal{X}, U)$ contains a neighbourhood of zero.*

In the car example, we can easily confirm the need for discontinuous feedbacks. Consider a neighbourhood of the origin with θ small; say $\mathcal{X} := \{(x, y, \theta) : x < 1,\ y < 1,\ \theta < \pi/6\}$. Then, its image by f for any possible control has no points of the form $(\dot{x}, \dot{y}, \dot{\theta})$ with $\dot{x} < 0$ (see Fig. 4). Therefore there is no open set containing the origin included in $f(\mathcal{X}, U)$, which confirms the need for discontinuous feedbacks. It can also be checked that the Brockett integrator and the unicycle require such feedbacks for stabilization; see [6, 39].

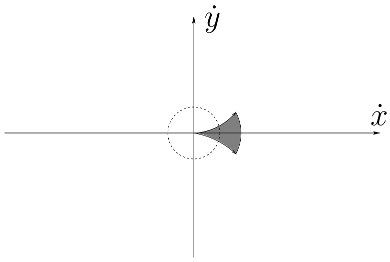

Figure 4: The image $f(x, y, \theta)$ for small values of θ projected in the $\dot{x}\dot{y}$ axis.

However, if we allow discontinuous feedbacks, it might not be clear what is the solution of the dynamic differential equation. The trajectory x, solution to

$$\dot{x}(t) = f(x(t), k(x(t))), \quad x(0) = x_0$$

is just defined, in a classical sense (known as Caratheodory solution), if k is assumed to be continuous. If k is discontinuous, a solution might even not exist! (See [10] for a further discussion of this issue).

Remark 3.2 *One might be tempted to think that the question of existence of a solution is just a mathematical detail, since if one uses a discontinuous feedback in a real plant certainly some trajectory will be followed. The point, however, is that using ordinary differential equations together with the classical definition of trajectory is no longer an appropriate model to describe the behaviour of the plant subject to discontinuous feedbacks. Besides, the possibility of "switching" at an infinite frequency brings about some theoretical and practical difficulties well-known to the sliding-mode and hybrid-systems research community.*

There is a known paradox that illustrates the problem of defining a trajectory subject to a discontinuous feedback. Consider a runner, a cyclist, and a dog departing from the same point, at the same time, with the same directions, and with velocities v, $2v$, and $3v$, *respectively. The runner and the cyclist keep pursuing the same direction, while the dog uses the following "discontinuous strategy": it runs forward until it reaches the cyclist, then it turns and runs backward until it reaches the runner; at this point, it turns forward and repeats the procedure.*

The question is: where is the dog after a certain time T? (See Fig. 5.)

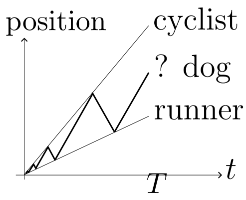

Figure 5: Undefined position of the dog.

The answer is: the dog can be **at any point** *between the cyclist and the runner.*

This might appear surprising since, apparently, every initial data is well defined and the final position of the dog is not. It is easy, however, to show that this is the answer if we reverse the time. Because, if the dog "starts" at time T from anywhere between the cyclist and the runner, going backward in time, they will all meet at the origin at time zero.

A possible way out of this undefined situation is to force the dog to maintain its strategy for a certain time $\delta > 0$ before deciding where it should turn to the next time. The dog trajectory would then be well-defined and unique (see Fig. 6). This idea of forcing a decision to be valid for a strictly positive time is used in the CLSS solution concept and in the MPC framework described.

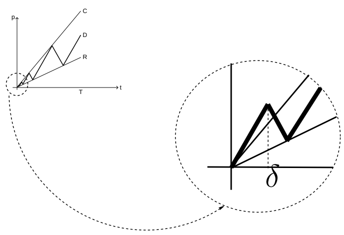

Figure 6: Well-defined position of the dog.

A question then arises of how to define a solution to a dynamic equation with a discontinuous feedback law. As is often the case for nontrivial questions, more than one definition has been proposed. The best known of such definitions is the concept of Filippov solution [16], which has a number of desirable properties. However, it was shown [37, 13], that there are controllable systems — nonholonomic systems, for example — that cannot be stabilized, even allowing discontinuous feedbacks, if the trajectories are interpreted in a Filippov sense.

A solution concept that has been proved successful in dealing with stabilization by discontinuous feedbacks for a general class of controllable systems is the concept of CLSS solution proposed by Clarke, Ledyaev, Sontag and Subbotin [11]. One of the main features of the concept is precisely the use of strictly positive inter-sampling times. We proceed to describe it. Consider a sequence of sampling instants $\pi := \{t_i\}_{i\geq 0}$ in $[0, +\infty)$ with $t_0 < t_1 < t_2 < \dots$ and such that $t_i \to \infty$ as $i \to \infty$. Let the function $t \mapsto \lfloor t \rfloor_\pi$ give the last sampling instant before t, that is

$$\lfloor t \rfloor_\pi := \max_i \{t_i \in \pi : t_i \leq t\}.$$

For a sequence π, the trajectory of the system under the feedback k is defined by

$$\dot{x}(t) = f(x(t), k(x(\lfloor t \rfloor_\pi))), \quad t \in [0, +\infty) \tag{1a}$$

$$x(t_0) = x_0. \tag{1b}$$

Here, the feedback is not a function of the state at every instant of time, rather it is a function of the state at the last sampling instant.

It is argued in a later section, that MPC can be combined naturally with a "sampling-feedback" law and thus define a trajectory in a way which is very similar to the concept introduced in [11]. Those trajectories are, under mild conditions, well-defined even when the feedback law is discontinuous.

4 Discontinuous Optimal Controls

Consider an MPC algorithm solving open-loop optimal control problems as the following.

$$\begin{aligned} &\text{Minimize} && \int_0^T L(x(s), u(s))ds + W(x(T)) \\ &\text{subject to} && \\ & && \dot{x}(s) = f(x(s), u(s)) \quad \text{a.e. } s \in [0, T] \\ & && x(t) = x_0 \\ & && u(s) \in U \quad \text{a.e. } s \in [0, T] \\ & && x(T) \in S. \end{aligned}$$

The data of this problem include an initial state vector $x_0 \in \mathbb{R}^n$, an input constraint set $U \subset \mathbb{R}^m$, a terminal state constraint set $S \in \mathbb{R}^n$, the running cost and terminal cost functions $L : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}$ and $W : \mathbb{R}^n \mapsto \mathbb{R}$, and the dynamics function $f : \mathbb{R}^n \times \mathbb{R}^m \mapsto \mathbb{R}^n$.

Let the pair trajectory/control $(\bar{x}, \bar{u})$ denote a minimizer for this problem (see e.g. [44] for a definition of a minimizer and a further discussion of optimal control).

A common assumption in most (continuous-time) MPC approaches is to impose that the control functions $\bar{u} : [0, T] \to \mathbb{R}^m$ obtained as solutions to the open-loop optimal control problems must be continuous, at least at the initial time. This assumption is very difficult to guarantee, specially when systems with constrained inputs are considered. It is well-known that even optimal control of linear systems frequently results in bang-bang control (i.e. control values are either at the maximum or the minimum limits). This arises, for example, if a minimum-time or a linear cost function is used; see e.g. [27].

At this point the following questions arise naturally: Why cannot the continuity assumption be simply relaxed? Why do most continuous-time MPC frameworks impose such continuity?

The difficulties of relaxing the continuity assumption result from the fact that most continuous-time MPC frameworks, after solving an optimal control problem at an instant of time, implement the value of the optimal control solution at that instant; i.e. implement the first value of the control function obtained. However, the value of the optimal control at an instant is not defined if a discontinuous control is allowed.

Suppose some control function $\bar{u}$ is a solution to an open-loop optimal control problem. Then, consider selecting a finite, or even countable, set of instants of time and change the value of the control at these instants of time to an arbitrary value. The control function obtained in this way is still optimal. This is because we are just changing the control values on a set of instants forming a set of zero Lebesgue measure. The values of f and L remain the same almost everywhere, thus the trajectory and the cost (which are obtained by integration) remain the same. It is now easy to see that using an MPC strategy that selects a single value of a control function to implement — which can be arbitrary if the control is allowed to be discontinuous — has serious problems. The resulting strategy is unpredictable, it can in fact be anything!

It is, of course, possible to impose that the optimal control problems search for a solution among continuous (or piecewise right-continuous) control functions only. This is the approach taken, for example, in [34, 35, 7, 8]. Such procedure, however raises some problems. It is true that we can approximate a discontinuous function by a continuous one with arbitrary precision. But, this does *not* imply that the minimum cost obtained using a continuous control, can be arbitrarily close to the optimal cost obtained allowing discontinuous controls. In fact, there are examples of problems for which the difference in the minimum costs between solutions of different regularity has a strictly positive gap that cannot be bridged. This phenomenon is known as the Lavrentiev phenomenon, and it might even arise in problems with very simple dynamics, such as $\dot{x} = u$; see [9, 44].

Another problem of imposing continuity is that most (if not all) results on the existence of a solution to a nonlinear optimal control problem consider problems where the admissible controls are selected among measurable functions [44]. This is significant because establishing existence of a (optimal) solution is important to guarantee stability of an MPC strategy [29, 19] unless a criterion requiring merely improvement in the cost of each solution is employed [38].

Other options to avoid the problem of discontinuous controls include using a discrete-time approach (e.g. [14]), or impose demanding hypotheses that guarantee continuity of the controls. For example, [31] impose a set of hypothesis implying that the value function is continuously differentiable. In [24] the same properties are achieved by considering just unconstrained inputs and unconstrained optimal control problems.

It is possible to use a continuous-time MPC framework that allows discontinuous controls, despite the difficulties mentioned. In [21, 20] the control that the MPC strategy implements is not an instant value of the control; rather it is a limiting behaviour of the optimal control. The control value u^* that is implemented at instant t is a control value satisfying

$$f(\bar{x}(t), u^*(t)) = \liminf_{\delta \downarrow 0} \frac{1}{\delta} \int_t^{t+\delta} f(\bar{x}(s), \bar{u}(s))\, ds.$$

Assuming continuity of f and convexity of the velocity set $f(\bar{x}, U)$, guarantees that u^* exists and is in U even when $\bar{u}$ is discontinuous.

A curious approach reported for systems with two inputs [3] is the use of piecewise semi-constant feedbacks. Each control is maintained equal to zero for half the sampling interval, and on the other half has a constant value dependent on the value of the state at the beginning of the interval. Selecting the control value to depend on the value of state at the beginning of the sampling interval is other technique to overcome the discontinuity problem.

Another alternative is to implement not a single control value, but to implement a portion of the optimal control obtained for an interval of nonzero length $[t_i, t_i + \delta)$, with $\delta > 0$:

$$u^*(t)) = \bar{u}(t), \quad t \in [t_i, t_i + \delta).$$

This is the approach taken, for example, in [34, 7, 8, 19], although [34, 7, 8] do require continuity. The work [28] imposes constant controls during the sampling interval which might reduce the computational burden of the optimization algorithms. Despite the fact that there are several works implementing the control during a strictly positive inter-sampling interval, this benefit is seldom exploited to relax the continuity assumption.

5 MPC Frameworks Allowing Discontinuity

There are in the literature a few works allowing discontinuous feedback laws in the context of MPC. The majority of these works consider discrete-time systems, avoiding this way most of the difficulties caused by discontinuous feedback laws; e.g. [33, 2]. The work [3] addresses continuous-time systems in power form by transforming them into discrete-time using the piecewise semi-constant controls described earlier, and then applying a discrete-time MPC algorithm. An early exception to the discrete-time case is the work of Michalska and Vinter in 1994 [35] that uses a continuous-time framework and in which the trajectories resulting from a discontinuous feedback law are interpreted as Filippov solutions. However, as was discussed previously, it was proved in the same year that it is not possible to stabilize the trajectories of some controllable nonlinear systems if these trajectories are interpreted in the sense of Filippov [37, 13].

Therefore, the class of nonlinear systems that could be addressed by continuous-time MPC with guaranteed stability was limited, not including the examples described in Section 2, among others. Despite that, there were a few works that appear to indicate that MPC was an appropriate framework to address problems of mobile robotics, for example. An illustration is the simulation results in [42, 15]. However, in the former work no stability analysis was carried out. In the latter it is recognized that MPC is a framework that could, in principle, address nonholonomic systems; but, their assumption on the stabilizability of the linearized system prevents them to guarantee stability.

Recently, MPC frameworks using a sampling feedback similar to the CLSS solution concept — which is capable of overcoming the difficulties caused by the use of discontinuous feedbacks and discontinuous optimal controls — was proposed [18, 19]. Even more recently there were other MPC works exploiting the advantages of the mentioned solution concept (see Alamir [1]).

The essential feature of those frameworks to allow discontinuities is simply the appropriate use of a positive inter-sampling time, combined with an appropriate interpretation of a solution to a discontinuous differential equation. Several other approaches proposed in the literature could, we believe, be easily adapted to handle discontinuities.

We proceed to describe some of the main ingredients of an MPC framework allowing discontinuities. (The time invariant case is considered for simplicity; see [19] for a framework allowing time-varying systems.)

Consider a plant with input constraints, where the evolution of the state after time t is predicted by the following nonlinear model.

$$\dot{x}(s) = f(x(s), u(s)) \quad \text{a.e. } s \geq t \tag{2a}$$

$$x(t) = x_t \tag{2b}$$

$$u(s) \in U. \tag{2c}$$

where the set U contains the origin and $f(0,0) = 0$. It is also possible to include state constraints.

$$x(t) \in X \subset \mathbb{R}^n \quad \text{for all } s \geq t.$$

The only difficulties added by the state constraints are in guaranteeing feasibility and in numerically solving the optimal control problems.

Consider also a sequence of sampling instants $\pi := \{t_i\}_{i\geq 0}$ with a constant inter-sampling time $\delta > 0$ (smaller than the horizon T) such that $t_{i+1} = t_i + \delta$ for all i. The feedback control is obtained by repeatedly solving online open-loop optimal control problems $\mathcal{P}(x_t, T)$ at each sampling instant $t \in \pi$, every time using the current measure of the state of the plant x_t.

$$\begin{aligned}
\mathcal{P}(x_t, T): \quad & \text{Minimize} \quad \int_t^{t+T} L(x(s), u(s))ds + W(x(t+T)) \\
& \text{subject to} \\
& \qquad \dot{x}(s) = f(x(s), u(s)) \quad \text{a.e. } s \in [t, t+T] \\
& \qquad x(t) = x_t \\
& \qquad x(s) \in X \quad \text{for all } s \in [t, t+T] \\
& \qquad u(s) \in U \quad \text{a.e. } s \in [t, t+T] \\
& \qquad x(t+T) \in S.
\end{aligned}$$

Here, we use the convention that the variable t represents real time while s denotes the time variable used in the prediction model. The vector x_t denotes the actual state of the plant measured at time t. The pair $(\bar{x}, \bar{u})$ denotes an optimal solution to an open-loop optimal control problem. The process (x^*, u^*) is the closed-loop trajectory and control resulting from the MPC strategy.

When discussing stability in the next section we will refer to the *design parameters*, which are the variables present in the open-loop optimal control problem that we are able to choose; these comprise the time horizon T, the running and terminal costs functions L and W, and the terminal constraint set $S \subset \mathbb{R}^n$.

The MPC algorithm consists of performing the following steps at a certain instant t_i.

1. Measure the current state of the plant x_{t_i}.

2. Compute the open-loop optimal control $\bar{u} : [t_i, t_i + T] \to \mathbb{R}^n$ solution to $\mathcal{P}(x_{t_i}, T)$.

3. Apply to the plant the control $u^*(t) := \bar{u}(t)$ in the interval $[t_i, t_i + \delta)$ (discard the remaining control $\bar{u}(t), t \geq t_i + \delta$).

4. Repeat the procedure from (1.) for the next sampling instant t_{i+1}.

The resultant control law is a "sampling-feedback" control since during each sampling interval, the control u^* is dependent on the state x_{t_i}. For each $i = 0, 1, \ldots$ let $\bar{u}^i$ be in the solution to $\mathcal{P}(x_{t_i}, T)$, then

$$x^*(t_i) = x_{t_i}, \quad \dot{x}^*(t) = f(x^*(t), k(t, x_{t_i})) \quad t \in [t_i, t_i + \delta),$$

where

$$k(t, x_{t_i}) = u^*(t) = \bar{u}^i(t) \quad t \in [t_i, t_i + \delta).$$

Equivalently

$$x^*(t_0) = x_{t_0}, \quad \dot{x}^*(t) = f(x^*(t), k(t, x^*(\lfloor t \rfloor_\pi))) \quad t \geq t_0,$$

where

$$k(t, x^*(\lfloor t \rfloor_\pi)) = u^*(t) \quad t \geq t_0.$$

This allows our MPC framework to overcome the inherent difficulty of defining solutions to differential equations with discontinuous feedbacks. In this way, the class of nonlinear systems potentially addressed by MPC is enlarged, including, for example, nonholonomic systems.

The convenient use of a strictly positive inter-sampling interval $\delta > 0$ has several advantages. In addition to being necessary for implementation of the MPC algorithm (since it is not conceivable to solve optimization problems at every possible instant of a continuous time interval), it allows discontinuous feedbacks and discontinuous optimal controls.

6 Stability

We start by describing a notion of stability required in the context of sampling-feedbacks and in the examples discussed.

Definition 6.1 The sampling-feedback k is said to *stabilize* the system (1) on X_0 if there exists a sufficiently small inter-sample time δ such that the following condition is satisfied. For any $\gamma > 0$ we can find a scalar $M > 0$ such that for any pair $(t_0, x_0) \in \mathbb{R} \times X_0$ we have $\|x(s + t_0; t_0, x_0, k)\| < \gamma$ for $s \geq M$.

Remark 6.1 *Note that this notion of stability establishes just the usual notion of attractiveness, saying that*

$$\|x(t)\| \to 0 \text{ as } t \to +\infty,$$

but does not imply the usual notion of Lyapunov stability

$$\forall \epsilon > 0 \, \exists \delta > 0 : \|x(0)\| < \delta \Rightarrow \|x(t)\| < \epsilon.$$

This latter notion follows automatically when a continuous differentiable Lyapunov function is employed to establish attractiveness.

The traditional definition of asymptotical stability in the literature requires both attractiveness and Lyapunov stability to be satisfied [23, 43]. An exception is the definition in [12] which, as the definition here, requires merely attractiveness.

Our choice is justified by the fact that attractiveness and Lyapunov stability are impossible to satisfy simultaneously for some systems that we would intuitively classify as controllable.

One such system is the car-like vehicle. Lyapunov stability requires that the state stays arbitrarily close to the origin, provided the initial state is close enough to the origin. However, even if we start arbitrarily close to the origin, we might have to move to a certain minimal distance away from the origin in order to drive to it. Select an $\epsilon > 0$ *smaller that* R_{min}*, and suppose the initial state is* $x = 0$*,* $y = \delta'$ *and* $\theta = 0$*. Even when selecting* δ' *arbitrarily small, the state must leave the* ϵ*-radius ball to drive to the origin. Therefore, to satisfy attractiveness, we cannot satisfy Lyapunov stability for this system.*

Several of the recent results on the stability of the MPC strategy use a terminal cost (W) as a major player to ensure stability.

It has been known for some time that using the terminal cost to be the integral of the running cost till infinity (i.e. $W(x(t+T)) = \int_{t+T}^{\infty} L(x(s), u(s))ds$) would make the optimization problem equivalent to an infinite horizon one and therefore would guarantee stability.

That infinity tail cost is hard to compute except in the linear case — where it coincides with the well-known quadratic Lyapunov function —, and in the nonlinear case when $x(t+T) = 0$ — the case in which $W(x(t+T))$ is trivially computed to be zero. Not surprisingly, the first results appeared for linear systems [5, 36] and also for nonlinear systems for which the terminal state $x(t+T)$ is constrained to the origin [25, 31]. Then, systems that could be considered almost as linear near to the origin started to be addressed [7, 14] through a development of the dual-mode approach [34]. These approaches required the terminal state $x(t+T)$ to be near to origin (i.e. the terminal set S was a ball/elipsoid around the origin), and after that a linear feedback could be used to compute the tail cost.

It is a known result by now, that similarly to the linear case, for the unconstrained problems ($S = \mathbb{R}^n$) stability is ensured also if the terminal cost is a Control Lyapunov function (CLF), forming with the running cost a control Lyapunov pair (W, L) [24]. To be a control Lyapunov pair the functions W and L must be positive definite, satisfy some properness properties, W should be continuously differentiable, and the following infinitesimal decrease property has to be satisfied. For all $x \in \mathbb{R}^n$ there is a control $u \in U$ satisfying

$$\nabla W(x) \cdot f(x, u) \leq -L(x, u).$$

But again, global control Lyapunov functions are hard to compute for the more interesting nonlinear systems. In addition, if a global CLF function is known for some system, then it is trivial to find a stabilizing feedback for such system (simply define $k(x) := \text{argmin}_u \{\nabla W(x) \cdot f(x, u) + L(x, u)\}$). The purpose of MPC is then limited to improve the performance of an already stabilizing the control strategy.

In the more interesting constrained case, it was found that if the infinitesimal decrease property holds only within the terminal set S and the terminal set is invariant, then the same stability conclusion is achieved. That is, stability is guaranteed if, in addition to some requirements on the design parameters (see [19] for details), the following stability condition is satisfied. For each $x_t \in S$, there is a control function $\tilde{u} : [t, t + \delta] \to U$ such that for all $s \in [t, t + \delta]$ we have

$$\nabla W(x(s)) \cdot f(x(s), \tilde{u}(s)) \leq -L(x(s), \tilde{u}(s)), \tag{SCa}$$

$$x(s; t, x_t, \tilde{u}) \in S \tag{SCb}$$

A similar condition can be seen, in different degrees of generality, in several works. In the approaches [7, 14] we can see condition (SCa) when $W(x) = x^T P x$, u is a linear feedback, and S is a level set of W (i.e. $S := \{x : W(x) \leq \alpha\}$), thereby ensuring invariance. In [22, 8] a general function W is used, u is a continuous function, and S is a level set of W. The survey [32] identified the conditions (SCa) and (SCb) as common ingredients in most stabilizing MPC schemes and showed them to be sufficient for stability in the discrete-time case. In [19] this condition appears as above, where it is combined with other conditions to ensure the existence of a solution to the optimal control problems. (To simplify exposition we are assuming here existence of (optimal) solutions to the dynamic optimization problems.)

It is interesting to note that the main trend has been to choose the set S as a consequence of the choice W: typically S is a level set of W. This should be the approach to follow when the system behaves almost as linear near to the origin or when the input constraints are the main difficulty to control the system. However, for some intrinsic nonlinear systems it might not be easy to drive the system to the origin from every possible direction in the state space. The nonholonomic systems fall within this category. In this case, choosing the set S to be the level set of W might not be the best option.

The alternative, explored in [19] and in the examples below, is first to identify a set of directions from which it is easy to drive the system to the origin, and then construct a function W satisfying the infinitesimal decrease condition (SCa).

There is one further generalization of condition (SCa) that can be easily done. That is to remove the differentiability assumption on W and substitute it by merely Lipschitz continuity. Using Rademacker's Theorem [12, 3.4.19] we may conclude that $s \mapsto W(x(s))$ is differentiable almost everywhere. The infinitesimal decrease condition then becomes the following. For each $x_t \in S$, there is a control function $\tilde{u} : [t, t + \delta] \to U$ such that for almost every $s \in [t, t + \delta]$ we have

$$\frac{d}{ds} W(x(s; t, x_t, \tilde{u})) \leq -L(x(s), \tilde{u}(s)), \tag{SCa'}$$

(We note that (SCa') is equivalent to (SCa) when W is continuously differentiable.) The proof of the stability results in [19] holds without modification since what is used in the proof of stability is the integral of (SCa) in the interval $[t, t + \delta]$ which remains unchanged.

This generalization of removing the differentiability assumption on W is advantageous for some systems. There are some systems that do not admit a smooth (continuously differentiable) CLF. The unicycle system studied here is precisely one of such systems. Because the unicycle system can be transformed into the nonholonomic integrator by an appropriate change of coordinates [39], and it has been shown by Artstein [4], that the nonholonomic integrator does not admit a smooth CLF.

For the use of CLF requiring even less regularity, with the infinitesimal decrease condition using Dini derivatives or proximal subgradients see for example [40, 11, 12].

Next, we analyse the selection of design parameters guaranteeing stability for the unicycle and car-like vehicle examples.

Example 1: The unicycle We proceed to analyse the unicycle system in Fig. 1 whose model is recalled here for convenience.

$$\begin{cases} \dot{x}(t) = (u_1(t) + u_2(t)) \cdot \cos\theta(t) \\ \dot{y}(t) = (u_1(t) + u_2(t)) \cdot \sin\theta(t) \\ \dot{\theta} = (u_1(t) - u_2(t)). \end{cases}$$

Here, the heading angle $\theta(t) \in [-\pi, \pi]$, and the controls $u_1, u_2(t) \in [-1, 1]$.

Our objective is to drive this system to the origin $(x, y, \theta) = (0, 0, 0)$.

Our task here is to find a set of design parameters that satisfies the stability conditions (SC) and therefore guarantee that the resulting MPC strategy is stabilizing.

One of the regions of the state space where it is easier to control this system is when the heading angle θ is pointing towards the origin of the plane $(x, y) = (0, 0)$. In this case, if the system just moves forward in a straight line (using $u_1 = u_2 = 1$) it is decreasing the x and y components of the state, while maintaining the heading angle θ. This set, with this strategy, is invariant provided we add the origin of the plane $(x, y) = (0, 0)$ for any heading angle θ.

It is convenient to define $\phi(x, y)$ to be the angle in $[-\pi, \pi]$ that points to the origin from position (x, y), that is

$$\phi(x,y) = \begin{cases} 0 & \text{if } x = 0, y = 0; \\ -(\pi/2)\,\mathrm{sign}(y) & \text{if } x = 0, y \neq 0; \\ \tan^{-1}(y/x) + \pi & \text{if } x > 0; \\ \tan^{-1}(y/x) & \text{if } x < 0. \end{cases}$$

(This function is easily implementable in Matlab, or other languages with a four-quadrant inverse tangent built-in function, as $\phi(x, y) = atan2(-y, -x)$.)

Now we are ready to define the design parameters. We define the terminal set S to be the set of states heading towards the origin of the plane together with the origin of the plane.

$$S := \{(x, y, \theta) \in \mathbb{R}^2 \times [-\pi, \pi] : \theta = \phi(x, y) \ \vee \ (x, y) = (0, 0)\}.$$

This set can be reached for any state if we define the horizon to be the time to complete an 180 degrees turn, that is $T = \pi/2$.

Define the running cost and the terminal cost functions simply as

$$L(x, y, \theta) = x^2 + y^2 + \theta^2, \qquad W(x, y, \theta) = 2\int_0^{\bar{t}} L(x(t), y(t), \theta(t))\, dt,$$

where $\bar{t}$ is the time to reach the origin. An explicit formula for the terminal cost is (see appendix) $W(x, y, \theta) = \frac{1}{3}(r^3 + |\theta|^3) + r\theta^2$ with $r = \sqrt{x^2 + y^2}$.

We proceed to check that condition (SCa') is satisfied, by analysing separately 3 subsets of S: **(i)** when $(x, y) = (0, 0)$; **(ii)** when $(x, y) \neq (0, 0)$ and $\theta = \phi(x, y) \neq 0$; and **(iii)** when $(x, y) \neq (0, 0)$ and $\theta = \phi(x, y) = 0$.

(i) Since $(x, y) = (0, 0)$, when $\theta = 0$ we are at the origin and so the control $u_1 = u_2 = 0$ satisfies (SCa'). In the remaining cases, when $\theta \neq 0$, the control is such that $u_1 = -u_2 = -\text{sign}(\theta)$. Then, the trajectory starting from $(0, 0, \theta_0)$ at $s = 0$ is

$$[x(s), y(s), \theta(s)] = [0, 0, \theta_0 - 2\,\text{sign}(\theta_0)s].$$

Therefore

$$W(x(s), y(s), \theta(s)) = \frac{1}{3}|\theta(s)|^3 = \frac{1}{3}|\theta_0 - 2\,\text{sign}(\theta_0)s|^3$$

and, when $\theta_0 > 0$ its time derivative at the point $s = 0$ is

$$\frac{d}{ds}W(x(s), y(s), \theta(s)) \;=\; \frac{d}{ds}\left[\frac{1}{3}[\theta_0 - 2s]^3\right]|_{s=0} \;=\; -2\theta^2 \;\leq\; -L(0, 0, \theta),$$

as required. The case when $\theta_0 < 0$ can be dealt with in a similar way.

(ii) When $\theta = \phi(x, y)$ and $(x, y) \neq (0, 0)$ the control strategy is to move forward, that is $u = (1, 1)$. The velocity is then $f(x, y, \theta, u) = (2\cos\theta, 2\sin\theta, 0) = (2\cos\phi, 2\sin\phi, 0)$. Using the fact that $\tan^{-1} z = \sin^{-1}\frac{z}{\sqrt{1+z^2}} = \cos^{-1}\frac{1}{\sqrt{1+z^2}}$ we have, in the case of $x > 0$,

$$\cos\phi = \cos[\tan^{-1}(y/x) + \pi] = -\cos[\tan^{-1}(y/x)] = \frac{-1}{\sqrt{1 + y^2/x^2}} = \frac{-x}{\sqrt{x^2 + y^2}}$$

$$\sin\phi = \sin[\tan^{-1}(y/x) + \pi] = -\sin[\tan^{-1}(y/x)] = \frac{-y/x}{\sqrt{1 + y^2/x^2}} = \frac{-y}{\sqrt{x^2 + y^2}}$$

Therefore $f(x, y, \theta, u) = (-2x/r, -2y/r, 0)$.

Computing the gradient of W we obtain $\nabla W(x, y, \theta) = (xr + 2x\theta^2/r, yr + 2y\theta^2/r, \theta^2\text{sign}(\theta) + 2\theta r)$, yielding

$$\nabla W(x, y, \theta) \cdot f(x, y, \theta, u) \;=\; -2(x^2 + y^2 + 2\theta^2) \;\leq\; -L(x, y, \theta)$$

as required. The cases when $x < 0$ and $x = 0$ can be verified in a similar way.

(iii) If $\theta_0 = \phi(x_0, y_0) = 0$ then $x_0 < 0$ and $y_0 = 0$. The trajectory starting at $(x_0, 0, 0)$ with controls $u_1 = u_2 = 1$ yields velocity $f(x, y, \theta, u) = (2, 0, 0)$ and trajectory $[x(s), y(s), \theta(s)] = [x_0 + 2s, 0, 0]$. Therefore

$$W(x(s), y(s), \theta(s)) = \frac{1}{3}|x(s)|^3 = -\frac{1}{3}(x_0 + 2s)^3$$

and, when $\theta_0 > 0$ its time derivative at the point $s = 0$ is

$$\frac{d}{ds}W(x(s), y(s), \theta(s)) \;=\; \frac{d}{ds}\left[-\frac{1}{3}[x_0 + 2s]^3\right]|_{s=0} \;=\; -2x_0^2 \;\leq\; -L(x, 0, 0),$$

as required.

Example 2: A car-like vehicle Recall the car-like vehicle described in Section 2.

$$\begin{cases} \dot{x} = v \cdot \cos\theta \\ \dot{y} = v \cdot \sin\theta \\ \dot{\theta} = v \cdot c \end{cases}$$

The control inputs v and c satisfy $v \in [0, v_{max}]$ and $c \in [-c_{max}, c_{max}]$. (Minimum turning radius $R_{min} = c_{max}^{-1}$.)

This example was already analysed in [19], but the main points are reproduced here since they are illustrative of the purpose of this paper. Our goal, as before, is to find a set of design parameters and verify that they satisfy the stability conditions (SCa) and (SCb).

To define S, a possibility is to look for feasible trajectories approaching the origin. One such set is represented in Fig. 7. This set is the union of all semi-circles with radius greater

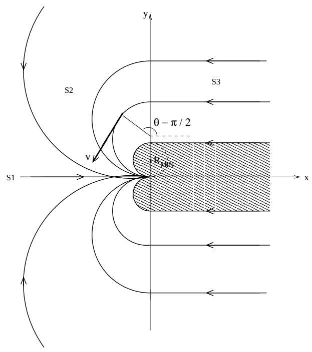

Figure 7: Set S of trajectories approaching the origin.

than or equal to R_{min}, with centre lying on the y axis, passing through the origin, and lying in the left half-plane. In order to make this set reachable in finite time from any point in the space, we add the set of trajectories that are horizontal lines of distance more than $2R_{min}$ from the x axis, and lie in the right half-plane. More precisely $S := S_1 \cup S_2 \cup S_3$ where

$$\begin{aligned} S_1 &:= \{(x, 0, 0) : x \leq 0\} \\ S_2 &:= \left\{(x, y, \theta) \in \mathbb{R}^2 \times [-\pi, \pi] : x \leq 0,\ x^2 + (y - r)^2 = r^2,\ r \geq R_{min},\ \tan\left(\theta - \frac{\pi}{2}\right) = \frac{y - r}{x}\right\} \\ S_3 &:= \{(x, y, \theta) \in \mathbb{R}^2 \times [-\pi, \pi] : x > 0,\ |y| \geq 2R_{min},\ \theta = \pi\}. \end{aligned}$$

The set S as defined above can be reached in a well-determined finite time-horizon. This horizon is the time to complete a circle of minimum radius at maximum velocity, that is $T = 2\pi R_{min}/v_{max}$. Choose L and W as

$$L(x, y, \theta) := x^2 + y^2 + \theta^2, \qquad W(x_0, y_0, \theta_0) := \int_0^{\bar{t}} L(x(t), y(t), \theta(t))dt,$$

where $\bar{t}$ is the time to reach the origin with the controls chosen to be the maximum velocity and the curvatures depicted in Fig. 7. We note that if we choose W to be just the sum of the squares of the components of the state, it would not work. This is because some components of the state actually increase along the trajectories in S_2.

It can be shown that this choice of design parameters satisfies our stability conditions. (A detailed verification of (SC) for each of the subsets S_1, S_2 and S_3 can be found in [17].) It follows that stability of the closed-loop trajectory is guaranteed.

7 Conclusions

A continuous-time MPC can be used with discontinuous optimal controls to generate stabilizing discontinuous feedback controls. Such MPC framework can stabilize a large class of nonlinear systems, including the nonholonomic systems which frequently appear in robotics and other areas of interest.

Some care is required to select the appropriate notion of stability, since for some nonlinear systems like the car-like vehicle the standard notion of asymptotic stability (requiring simultaneously attractiveness and Lyapunov stability) is impossible to achieve.

The use of a strictly positive inter-sampling time in a continuous-time MPC framework is essential when discontinuous feedback controls or discontinuous optimal control are to be employed. In addition, such use is much more natural since it allows some time to solve the optimization problems on-line.

A rigorous stability analysis can be carried out for an MPC framework possessing the above features.

Acknowledgement

The financial support from project FCT POSI/EEI/14156/1998 and from CMAT under programme POCTI of QCAIII is gratefully acknowledged.

References

[1] M. Alamir. New path-generation based receding horizon formulation for constrained stabilisation of nonlinear systems. In *40th IEEE Conference on Decision and Control*, pages 496–501, Orlando, Florida, 2001.

[2] M. Alamir and G. Bornard. Globally stabilising discontinuous feedback law for a class of nonlinear systems. *European Journal of Control*, 2(1):44–50, 1996.

[3] M. Alamir and H. Khennouf. Discontinuous receding horizon based stabilizing feedback for nonholonomic systems in power form. In *34th IEEE Conference on Decision and Control*, pages 4300–4, New Orleans, LA, 1995.

[4] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Analysis, Theory, Methods and Applications*, 7(11):1163–1173, 1983.

[5] R. R. Bitmead, M. Gevers, and V. Wertz. *Adaptive Optimal Control: The Thinking Man's GPC*. Prentice Hall, New York, 1990.

[6] R. W. Brockett. Asymptotic stability and feedback stabilization. In R. W. Brockett, R. S. Millman, and H. S. Sussmann, editors, *Differential Geometric Control Theory*, pages 181–191. Birkhouser, Boston, 1983.

[7] H. Chen and F. Allgöwer. A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 34(10), 1998.

[8] W.H. Chen, D. J. Ballace, and J. O'Reilly. Model predictive control of nonlinear systems: computational delay and stability. *IEE Proceedings, Part D: Control Theory and Applications*, 147:387–394, 2000.

[9] F. H. Clarke. *Methods of Dynamic and Nonsmooth Optimization.* SIAM, 1989.

[10] F. H. Clarke. Nonsmooth analysis in control theory: a survey. *European Journal of Control; Special issue: Fundamental Issues in Control*, 7:145–159, 2001.

[11] F. H. Clarke, Y. S. Ledyaev, E. D. Sontag, and A. I. Subbotin. Asymptotic controllability implies feedback stabilization. *IEEE Transactions on Automatic Control*, 42(10):1394–1407, 1997.

[12] F. H. Clarke, Y. S. Ledyaev, R. J. Stern, and P. R. Wolenski. *Nonsmooth Analysis and Control Theory.* Springer Verlag, New York, 1998.

[13] J. M. Coron and L. Rosier. A relation between time-varying and discontinuous feedback stabilization. *Journal of Math. Systems, Estimation and Control*, 4:67–84, 1994.

[14] G. De Nicolao, L. Magni, and R. Scattolini. Stabilizing receding-horizon control of nonlinear time-varying systems. *IEEE Transactions on Automatic Control*, 43(7):1030–1036, 1998.

[15] S. L. de Oliveira Kothare and M. Morari. Contractive model predictive control for constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 45:1053–1071, 2000.

[16] A. F. Filippov. *Differential Equations with Discontinuous Righthand Sides.* Kluwer Academic, Dordrecht, 1988.

[17] F. A. C. C. Fontes. *Optimisation-Based Control of Constrained Nonlinear Systems.* PhD thesis, Imperial College of Science Technology and Medicine, University of London, London SW7 2BY, U.K., 1999.

[18] F. A. C. C. Fontes. Discontinuous feedback stabilization using nonlinear model predictive controllers. In *Proceedings of CDC 2000 – 39th IEEE Conference on Decision and Control,*, Sydney, Australia, December 2000.

[19] F. A. C. C. Fontes. A general framework to design stabilizing nonlinear model predictive controllers. *Systems & Control Letters*, 42:127–143, 2001.

[20] F. A. C. C. Fontes and R. B. Vinter. Design of nonlinear model predictive controllers with specified rates of exponential decay. Technical Report C00.01, Centre for Process Systems Engineering, Imperial College, London SW7 2BY, U.K., 2000.

[21] F. A. C. C. Fontes and R. B. Vinter. Nonlinear model predictive control: Specifying rates of exponential stability without terminal state constraints. In *Proceedings of ADCHEM 2000 - IFAC Meeting on Advanced Control of Chemical Processes*, Pisa, Italy, June 2000.

[22] E. Gyurkovics. Receding horizon control via Bolza-time optimization. *Systems and Control Letters*, 35:195–200, 1998.

[23] W. Hahn. *Stability of motion.* Springer-Verlag, New York, 67.

[24] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46:776–782, 2001.

[25] S. Keerthi and E. Gilbert. Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon aproximations. *Journal of Optimization Theory and Applications*, 57:265–293, 1988.

[26] I. Kolmanovsky and N.H. McClamroch. Developments in nonholonomic control problems. *IEEE Control Systems*, pages 20–36, December 1995.

[27] J. Macky and A. Strauss. *Introduction to Optimal Control Theory.* Springer-Verlag, Berlin, 1982.

[28] L. Magni, R. Scattolini, and K. Aström. Global stabilization of the inverted pendulum using model predictive control. In *Proceedings of the IFAC World Congress 2002*, Barcelona, Spain, 2002.

[29] D. Q. Mayne. Nonlinear model predictive control: An assessment. In *5th International Conference on Chemical Process Control*, volume 93 of *AIChe Symposium Series*, pages 217–231, 1997.

[30] D. Q. Mayne and H. Michalska. An implementable receding horizon controller for stabilization of nonlinear systems. In *Proceedings of the 29th IEEE Conference on Decision and Control*, pages 3396–7, Honolulu, Hawaii, 1990.

[31] D. Q. Mayne and H. Michalska. Receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 35:814–824, 1990.

[32] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.

[33] E. S. Meadows, M. A. Henson, J. E. Eaton, and J. B. Rawlings. Receding horizon control and discontinuous state feedback stabilization. *International Journal of Control*, 62(5):1217–1229, 1995.

[34] H. Michalska and D. Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 38:1623–1633, 1993.

[35] H. Michalska and R. B. Vinter. Nonlinear stabilization using discontinuous moving-horizon control. *IMA Journal of Mathematical Control and Information*, 11:321–340, 1994.

[36] J. B. Rawlings and K. R. Muske. The stability of constrained receding horizon control. *IEEE Transactions on Automatic Control*, 38(10):1512–1516, 1993.

[37] E. P. Ryan. On Brockett's condition for smooth stabilizability and its necessity in a context of nonsmooth feedback. *SIAM Journal of Control and Optimization*, 32:1597–1604, 1994.

[38] P. O. M. Scokaert, D. Q. Mayne, and J. B. Rawlings. Suboptimal model predictive control (feasibility implies stability). *IEEE Transactions on Automatic Control*, 44(3):648–654, 1999.

[39] E. Sontag. Stability and stabilization: discontinuities and the effect of disturbances. In F. H. Clarke and R. Stern, editors, *Nonlinear Analysis, Differential Equations and Control*, pages 551–598. Kluwer, Dordrecht, 1998.

[40] E. D. Sontag. A Lyapunov-like characterization of asymptotic controllability. *SIAM Journal of Control and Optimization*, 21:462–471, 1983.

[41] E. D. Sontag and H. J. Sussman. Remarks on continuous feedback. In *IEEE Conference on Decision and Control, Albuquerque*, pages 916–921, 1980.

[42] H. Van Essen and H. Nijmeijer. Non-linear model predictive control of constrained mobile robots. In J. L. Martins de Carvalho, F. A. C. C. Fontes, and M. d. R. de Pinho, editors, *Proceedings of the European Control Conference - ECC'01*, pages 1157–1162, Porto, Portugal, 2001.

[43] M. Vidyasagar. *Nonlinear Systems Analysis.* Prentice Hall, Englewood Cliffs, NJ, 1993.

[44] R. B. Vinter. *Optimal Control.* Birkhauser, Boston, 2000.

Appendix

Here we compute an explicit formula for the terminal cost $W(x_0, y_0, \theta_0) = 2\int_0^{\bar{t}} L(x(t), y(t), \theta(t))\, dt$, where $\bar{t}$ is the time to reach the origin from any point in S. We compute separately the cost $W1$ of reaching the origin of the state space when we are already at the origin of the plane, and $W2$ the cost to reach the origin of the plane when $\theta_0 = \phi(x_0, y_0)$. The cost W is then defined outside S by extending to the whole space the formula obtained.

(i) When $(x_0, y_0) = 0$, and the controls are $u_1 = -u_2 = -\text{sign}\,\theta_0$. The trajectory is $[x(t), y(t), \theta(t)] = [0, 0, \theta_0 - 2\,\text{sign}(\theta_0)t]$ and the time to reach the origin is $t' = |\theta_0|/2$. Therefore

$$W1 = 2\int_0^{|\theta_0|/2} L(x(t), y(t), \theta(t))\,dt = 2\int_0^{|\theta_0|/2} (\theta_0 - 2\text{sign}\,\theta_0 t)^2\,dt = \frac{1}{3}|\theta_0|^3.$$

(ii) When $\theta_0 = \phi(x_0, y_0)$, the controls are $u_1 = u_2 = 1$. The velocity is $[\dot{x}(t), \dot{y}(t), \dot{\theta}] = [2 \cdot \cos\phi, 2 \cdot \sin\phi, 0]$. In polar coordinates (r, ψ) that is $[\dot{r}(t), \dot{\psi}(t), \dot{\theta}] = [-2, 0, 0]$, yielding the trajectory $[r(t), \psi(t), \theta(t)] = [r_0 - 2t, \psi_0, \theta_0]$ and the time to reach the origin of the plane is $t'' = r_0/2$. Therefore

$$W2 = 2\int_0^{r_0/2} L(x(t), y(t), \theta(t))\,dt = 2\int_0^{r_0/2} (r^2(t) + \theta^2(t))\,dt = \frac{1}{3}r_0^3 + r_0\theta_0^2.$$

Finally, computing the total path to the origin we obtain

$$W(x_0, y_0, \theta_0) = 2\int_0^{\bar{t}} L(x(t), y(t), \theta(t))\,dt = W1 + W2 = \frac{1}{3}r_0^3 + \frac{1}{3}|\theta_0|^3 + r_0\theta_0^2.$$