Hybrid MPC: Open-Minded but Not Easily Swayed

S. Emre Tuna, Ricardo G. Sanfelice, Michael J. Messina, and Andrew R. Teel

Department of Electrical and Computer Engineering, University of California, Santa Barbara, CA 93106, USA
{emre,rsanfelice,mmessina,teel}@ece.ucsb.edu

Summary. The robustness of asymptotic stability with respect to measurement noise for discrete-time feedback control systems is discussed. It is observed that, when attempting to achieve obstacle avoidance or regulation to a disconnected set of points for a continuous-time system using sample and hold state feedback, the noise robustness margin necessarily vanishes with the sampling period. With this in mind, we propose two modifications to standard model predictive control (MPC) to enhance robustness to measurement noise. The modifications involve the addition of dynamical states that make large jumps. Thus, they have a hybrid flavor. The proposed algorithms are well suited for the situation where one wants to use a control algorithm that responds quickly to large changes in operating conditions and is not easily confused by moderately large measurement noise and similar disturbances.

1 Introduction

1.1 Objectives

The first objective of this paper is to discuss the robustness of asymptotic stability to measurement noise for discrete-time feedback control systems. We focus on control systems that perform tasks such as obstacle avoidance and regulation to a disconnected set of points. We will compare the robustness induced by pure state feedback algorithms to the robustness induced by *dynamic* state feedback algorithms that have a "hybrid" flavor. Nonlinear model predictive control (MPC), in its standard manifestation, will fall under our purview since 1) it is a method for generating a pure state feedback control (see [14] for an excellent survey), 2) it can be used for obstacle avoidance (see [11, 12, 18]) and regulation to a disconnected set of points (this level of generality is addressed in [9] for example), and 3) dynamic "hybrid" aspects can be incorporated to enhance robustness to measurement noise. The second objective of this paper is to demonstrate such hybrid modifications to MPC. The proposed feedback algorithms are able to respond rapidly to significant changes in operating conditions without getting confused by moderately large measurement noise and related disturbances. The findings in this paper are preliminary: we present two different hybrid modifications of MPC, but we have not investigated sufficiently the differences between these modifications, nor have we characterized their drawbacks.

R. Findeisen et al. (Eds.): Assessment and Future Directions, LNCIS 358, pp. 17–34, 2007.
springerlink.com

1.2 What Do We Mean by "Hybrid MPC"?

First we discuss the term "hybrid" and consider how it has appeared before in the context of MPC.

"Hybrid" Dynamical Systems

In the context of dynamical systems, "hybrid" usually indicates systems that combine continuous and discrete aspects. Often "continuous" and "discrete" refer to the time domains on which solutions are defined. See, for example, [5, 13, 22]. In this situation, a hybrid dynamical system is one in which solutions are defined on time domains that combine continuous evolution and discrete evolution. (The time domain thus may be a subset of the product of the nonnegative reals and the nonnegative integers. See, for example, [3] and [5, 6, 7, 20]). The state flows continuously, typically via a differential equation, as hybrid time advances continuously; the state jumps, according to a update map or "difference" equation, as the hybrid time advances discretely. Whether flowing or jumping occurs depends on the state. The state may or may not contain logic variables that take values in a discrete set. If such variables exist, they do not change during the continuous evolution. Similarly, state variables that must evolve continuously do not change during jumps. We note here that a continuous-time control system implemented with a sample and hold device is a hybrid dynamical system of this type. Thus, when MPC based on a discrete-time model of a continuous-time process is used to synthesize a state feedback that is implemented with sample and hold, this can be thought of as hybrid control, although perhaps not as "hybrid MPC".

Other times, "continuous" and "discrete" refer to the domains in which the state components take values, while the time domain is fixed to be discrete. In other words, a hybrid system sometimes means a discrete-time system in which some of the variables can take on any of a continuum of values while other states can take on any values in a discrete set. This appears to be the most common meaning of "hybrid dynamical system" as used in the MPC literature, and we will mention specific work below.

"Hybrid" MPC

We believe that the development of MPC for hybrid systems that involve both flowing and jumping will be a very stimulating area of research. Nevertheless, throughout this paper, we will only consider discrete-time systems (although they can be thought of as coming from sampled continuous-time systems). Thus, our meaning of "hybrid MPC" must be related to the second one given above. The main feature of the MPC that we propose is that it is dynamic, sometimes introducing variables that take discrete values, with the aim of enhancing robustness to measurement noise. We focus on discrete-time control problems that can be solved robustly using pure state feedback but that can be solved more robustly by adding dynamics, perhaps with variables that take on discrete values. The idea of adding dynamics to improve robustness is not new, especially as it pertains to the control of continuous-time systems. See [19], [21], [2], [16]. Our purpose is to emphasize this observation in discrete time and to present general dynamic or "hybrid" algorithms that have potential for wide applicability, are simple conceptually, and that improve robustness to measurement noise.

Regarding results on hybrid MPC that have appeared in the literature previously, it is tempting to try to make distinctions between hybrid MPC for nonhybrid systems and (nonhybrid) MPC for hybrid systems. However, the distinction can be blurred easily by treating logic variables from the controller as part of the state of the plant to be controlled. The class of discrete-time hybrid systems to which MPC is most often applied is the class of so-called piecewise affine (PWA) control systems. The equivalence of this type of hybrid model to several other classes of hybrid models has been established in [10]. Model predictive control for PWA systems has been discussed in [1], where the optimization problems to generate the MPC feedback law are shown to be mixed integer multiparameter programs. In other work, the authors of [16] propose a hybrid MPC strategy for switching between a predetermined robust stabilizing state feedback controller and an MPC controller aimed at performance. An early result in [17] used a "dual-mode" approach that involved switching between MPC and a local controller.

2 Control Systems and Measurement Noise

2.1 Introduction

We consider the analysis and design of control algorithms for discrete-time systems of the form

$$x^+ = f(x, u) \ , \tag{1}$$

where $x \in \mathbb{R}^n$ denotes the *state*, x^+ the next value of the state, and $u \in \mathcal{U}$ the *control input*. The function f is assumed to be continuous. At times we will re-write the system (1) as $x^+ = x + \widetilde{f}(x,u)$, where $\widetilde{f}(x,u) := f(x,u) - x$, to emphasize that the discrete-time control system may represent the sampling of a continuous-time control system and that the next state value is not too far from the current state value, i.e., $\widetilde{f}$ is not very large.

In this paper we consider two types of feedback control algorithms: 1) pure state feedback, i.e., $u = \kappa_{\text{ßPSF}}(x)$, where $\kappa_{\text{ßPSF}} : \mathbb{R}^n \to \mathcal{U}$ is not necessarily continuous, and 2) dynamic state feedback, i.e., $u = \kappa_{\text{ßDSF}}(x, \xi)$, $\xi^+ = g_{\text{ßDSF}}(x, \xi)$, where $\xi \in \mathbb{N}$ and $\kappa_{\text{ßDSF}} : \mathbb{R}^n \times \mathbb{N} \to \mathcal{U}$ and $g_{\text{ßDSF}} : \mathbb{R}^n \times \mathbb{N} \to \mathbb{N}$ are not necessarily continuous, where $\mathbb{N} := \{0, 1, \ldots\}$. We are especially interested in the effect of measurement noise. In the case of pure state feedback, this means that $u = \kappa_1(x + e)$, where e represents measurement noise. In the case of dynamic state feedback, this means $u = \kappa_2(x + e, \xi)$, $\xi^+ = g(x + e, \xi)$. We focus on control problems where pure state feedback will have small measurement noise robustness margins, regardless of the control algorithm used. We will show below that problems of this type include controlling continuous-time systems using small sampling periods while attempting to achieve obstacle avoidance and/or regulation to a disconnected set of points.

We frame the discussion around three prototypical control tasks for the continuous-time control system $\dot{x} = v$, where $x \in \mathbb{R}^2$ and $v \in \overline{\mathbb{B}} \subset \mathbb{R}^2$ ($\overline{\mathbb{B}}$ denotes the closed unit ball and $\delta\overline{\mathbb{B}}$ denotes the closed ball of radius δ). The problems are to use sample and hold control with a relatively small sampling period to achieve 1) global regulation to a point, 2) global regulation to a set consisting of two (distinct) points, and 3) global regulation to a target while avoiding an obstacle. In each case, the discrete-time control system is $x^+ = x + u$, where $u \in \delta\overline{\mathbb{B}}$ and $\delta > 0$ represents the sampling period.

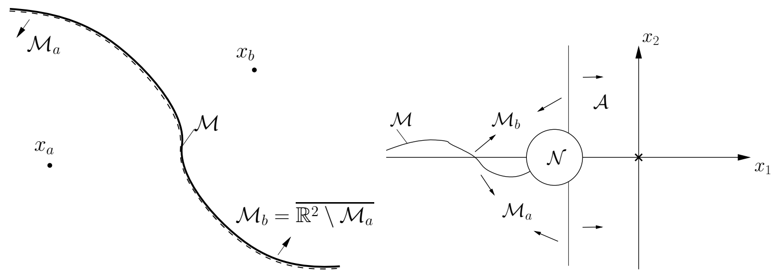

(a) Regulation to two distinct points. (b) Regulation to a target with obstacle avoidance.

Fig. 1. Global regulation to attractors

2.2 Global Regulation to a Point

Suppose we have designed a (family of) continuous feedback(s) $\kappa_\delta : \mathbb{R}^2 \to \delta\overline{\mathbb{B}}$ to achieve stability of and global asymptotic convergence to a point x^* which we take to be the origin without loss of generality. For example, suppose $\delta \in (0, 1]$ and we take

$$\kappa_\delta(x) = \frac{-\delta x}{\max\{1, |x|\}} \ . \tag{2}$$

To analyze the behavior of the system with measurement noise, define, for each $(s, \delta) \in \mathbb{R}_{\geq 0} \times (0, 1]$,

$$\gamma(s, \delta) = \frac{\max\{1, s\} - \delta}{\max\{1, s\}} \ . \tag{3}$$

Note that $\gamma(s, \delta) < 1$ for all $(s, \delta) \in \mathbb{R}_{\geq 0} \times (0, 1]$ and $\gamma(\cdot, \delta)$ is nondecreasing. Then note that

$$|x + \kappa_\delta(x + e)| \leq \frac{|x|\left(\max\{1, |x + e|\} - \delta\right) + \delta|e|}{\max\{1, |x + e|\}} \tag{4}$$

so that when $|e| \leq 0.5|x|$ we have $|x + \kappa_\delta(x + e)| \leq |x|\gamma(1.5|x|, 0.5\delta)$. It follows that the state trajectory will converge to a neighborhood of the origin that is proportional to the worst case size of the measurement noise, regardless of how small δ is. In other words, fast sampling does not make the system more and more sensitive to measurement noise.

2.3 Global Regulation to a Set Consisting of Two Distinct Points

Suppose we have designed a pure state feedback control algorithm $\kappa : \mathbb{R}^2 \to \delta\overline{\mathbb{B}}$ to achieve stability of and global asymptotic convergence to the set $\mathcal{A} := \{x_a, x_b\}$, where $x_a \neq x_b$. Stability implies that if the system starts close to one of the points it will stay close to that point forever. Let $\mathcal{H}_a$, respectively $\mathcal{H}_b$, denote the set of points that produce trajectories converging to x_a, respectively x_b. An illustration is given in Figure 1(a). By uniqueness of solutions, these sets are well defined and disjoint. By global asymptotic stability they cover $\mathbb{R}^2$, and because of the stability property each set is nonempty. We define $\mathcal{H}$ to be the intersection of the closures of $\mathcal{H}_a$ and $\mathcal{H}_b$, i.e., $\mathcal{H} = \overline{\mathcal{H}_a} \cap \overline{\mathcal{H}_b}$. Note that for each point $x \in \mathcal{H}$ there exists a neighborhood of x intersecting both $\mathcal{H}_a$ and $\mathcal{H}_b$. Again, using stability, it follows that $\mathcal{H}$ does not include neighborhoods of $\mathcal{A}$. Now, due to the nature of dynamical systems, the sets $\mathcal{H}_a$ and $\mathcal{H}_b$ are forward invariant. In particular

$$z_a \in \mathcal{H}_a \ , \ z_b \in \mathcal{H}_b \qquad \Longrightarrow \qquad z_a + \kappa(z_a) \in \mathcal{H}_a \ , \ z_b + \kappa(z_b) \in \mathcal{H}_b \ . \tag{5}$$

It can be shown that this fact has the following consequence (in the statement below, $\mathcal{H} + \delta\mathbb{B}$ denotes the set of points having distance less than δ from $\mathcal{H}$):

If $x \in \mathcal{H} + \delta\mathbb{B}$ then there exists e with $|e| < \delta$ such that $x + \kappa(x + e) \in \mathcal{H} + \delta\mathbb{B}$.

In turn it follows that, for each initial condition $x \in \mathcal{H}+\delta\mathbb{B}$ there exists a noise sequence $\mathbf{e} := \{e_k\}_{k=0}^{\infty}$ such that $|e_k| < \delta$ for all k and such that $\phi(k, x, \mathbf{e}) \in \mathcal{H} + \delta\mathbb{B}$ for all k, where $\phi(k, x, \mathbf{e})$ denotes the trajectory starting from x at the kth step under the influence of the measurement noise sequence $\mathbf{e}$. This (small when δ is small) noise sequence does not allow the trajectory to approach the attractor $\mathcal{A}$.

The above statement has the following explanation: Without loss of generality, suppose $x \in \mathcal{H}_a$. Since $x \in \mathcal{H} + \delta\mathbb{B}$, there exists z such that $|x - z| < \delta$ and $z \in \mathcal{H}_b$. In particular, $z + \kappa(z) \in \mathcal{H}_b$. Pick $e = z - x$ and consider $x + \kappa(x + e)$. If $x + \kappa(x + e) \in \mathcal{H}_b$ then $x + \kappa(x + e) \in \mathcal{H} + \delta\mathbb{B}$ since there must be a point on the line connecting x to $x+\kappa(x+e)$ that belongs to $\mathcal{H}$ and since the length of this line is less than δ since $|\kappa(x+e)| < \delta$. If $x+\kappa(x+e) \in \mathcal{H}_a$ then $x+\kappa(x+e) \in \mathcal{H}+\delta\mathbb{B}$ since there must be a point on the line connecting $x+\kappa(x+e)$ and $x+e+\kappa(x+e)$ that belongs to $\mathcal{H}$ and the length of this line must be less than δ since $|e| < \delta$.

To summarize, no matter how we build our pure state feedback algorithm, when δ is small there will be small noise sequences that can keep the system from converging toward the attractor.

2.4 Global Regulation to a Target with Obstacle Avoidance

Suppose we have designed a pure state feedback control algorithm $\kappa : \mathbb{R}^2 \backslash \mathcal{N} \to \delta\overline{\mathbb{B}}$ to achieve stability of and "global" convergence to a point x^* while avoiding an obstacle covering the set $\mathcal{N}$. The situation is depicted in Figure 1(b). For simplicity, we assume that the set $\mathcal{A}$ is made stable and attractive. Basically, this says that if the control finds that the vehicle is nearly past the obstacle it moves in the direction of the target. Let $\mathcal{H}_a$, respectively $\mathcal{H}_b$, denote the set of points that produce trajectories converging to $\mathcal{A}$ by crossing into the set $\mathcal{A}$ above the obstacle, respectively below the obstacle. By uniqueness of solutions, these sets are well defined and disjoint. By "global" asymptotic convergence they cover $\mathbb{R}^2 \backslash (\mathcal{N} \cup \mathcal{A})$, and because of stability and attractivity of $\mathcal{A}$ each set is nonempty. We define $\mathcal{H}$ to be the intersection of the closures of $\mathcal{H}_a$ and $\mathcal{H}_b$. Again using stability of $\mathcal{A}$ it follows that $\mathcal{H}$ does not include neighborhoods of $\mathcal{A}$. Because of this (5) holds, at least when δ is small enough. Using the same reasoning as in the previous subsection, we conclude that measurement noise of size δ can be used to keep the trajectories close to $\mathcal{H}$, which is on the "wrong" side of the obstacle, or else make the vehicle crash into the obstacle.

2.5 A General Principle

We point out here that the ideas put forth above in the discussion about stabilization of an attractor consisting of two distinct points and the discussion about obstacle avoidance generalize. Indeed, let $\mathcal{O} \subset \mathbb{R}^n$ be open and consider the discrete-time system

$$x^+ = x + \widetilde{f}(x) \ . \tag{6}$$

Let $\bar{h} \in \mathbb{N}_{\geq 2}$ and let the sets $\mathcal{H}_i$, for $i \in \{1, \ldots, \bar{h}\}$, satisfy $\bigcup_i \mathcal{H}_i = \mathcal{O}$. Define $\mathcal{H} = \bigcup_{i,j,i \neq j} \overline{\mathcal{H}}_i \cap \overline{\mathcal{H}}_j$.

Lemma 1. *Suppose that for each $z \in \mathcal{H}$ there exist $i, j \in \left\{1, \ldots, \bar{h}\right\}$ with $i \neq j$ and for each $\rho > 0$ there exist points $z_i, z_j \in \{z\} + \rho\mathbb{B}$ so that $z_i + \widetilde{f}(z_i) \in \mathcal{H}_i$ and $z_j + \widetilde{f}(z_j) \in \mathcal{H}_j$. Let $\varepsilon > 0$. If $x \in \mathcal{H} + \varepsilon\mathbb{B}$, $\{x\} + 2\varepsilon\mathbb{B} \subset \mathcal{O}$, and $|\widetilde{f}(x + e)| < \varepsilon$ for all $|e| < \varepsilon$ then there exists e such that $|e| < \varepsilon$ and $x + \widetilde{f}(x + e) \in \mathcal{H} + \varepsilon\mathbb{B}$.*

In turn, we have the following result.

Corollary 1. *Let $\varepsilon > 0$. Let $\mathcal{C} \subset \mathcal{O}$ be such that, for each $\xi \in \mathcal{C}$, $\xi + 2\varepsilon\mathbb{B} \subset \mathcal{O}$ and $|\widetilde{f}(\xi + e)| < \varepsilon$ for all $|e| < \varepsilon$. Then, for each $x_0 \in \mathcal{C} \cap (\mathcal{H} + \varepsilon\mathbb{B})$ there exists a sequence $\{e_k\}$ with $|e_k| < \varepsilon$ such that the sequence generated by $x_{k+1} = x_k + \widetilde{f}(x_k + e_k)$ satisfies $x_k \in \mathcal{H} + \varepsilon\mathbb{B}$ for all k such that $x_i \in \mathcal{C}$ for all $i \in \{0, \ldots, k-1\}$.*

A similar result applies to systems of the form $x^+ = x + \widetilde{f}(x, \kappa(x + e))$ and long as $\widetilde{f}(\cdot, u)$ is locally Lipschitz uniformly over u's in the range of κ. The ideas used to establish a result for such systems parallels the main idea in the proof of [4, Proposition 1.4]. We omit this result because of space limitations.

3 Standard MPC

In this section we review "standard MPC". In standard MPC a pure state feedback function is generated as the mapping from the state x to the solution to an optimization problem, parametrized by x, that uses continuous functions and does not use hard constraints. It has been shown in [15] that standard MPC yields a closed loop with some robustness to measurement noise. (This is in contrast to the situation where the MPC optimization involves hard constraints. Examples have been given in [8] to show that hard constraints can lead to zero robustness margins.) However, as suggested by the discussion in the previous section, the robustness margins may be quite small, especially if the discrete-time plant is coming from a discrete-time model of a continuous-time system using a relatively small sampling period and the control task is obstacle avoidance or regulation to a disconnected set of points.

The control objective is to keep the state in the open state space $\mathbb{X} \subset \mathbb{R}^n$ and stabilize the closed *attractor* $\mathcal{A} \subset \mathbb{X}$. MPC can be used to achieve this objective. The MPC algorithm is described as follows:

We denote an input sequence $\{u_0, u_1, \ldots\}$ by $\mathbf{u}$ where $u_i \in \mathcal{U}$ for all $i \in \mathbb{N}$. Let $\mathbb{E}_{\geq 0}$ denote $[0, \infty]$. Let $\sigma : \mathbb{R}^n \to \mathbb{E}_{\geq 0}$ be a *state measure* with the following properties: (i) $\sigma(x) = 0$ for $x \in \mathcal{A}$, $\sigma(x) \in (0, \infty)$ for $x \in \mathbb{X} \backslash \mathcal{A}$, and $\sigma(x) = \infty$ for $x \in \mathbb{R}^n \setminus \mathbb{X}$, (ii) continuous on $\mathbb{X}$, (iii) $\sigma(x)$ blows up as either x gets unbounded or approaches to the border of $\mathbb{X}$. We let $\ell : \mathbb{R}^n \times \mathcal{U} \to \mathbb{E}_{\geq 0}$ be the *stage cost* satisfying $\ell(x, u) \geq \sigma(x)$ and $g : \mathbb{R}^n \to \mathbb{E}_{\geq 0}$ the *terminal cost* satisfying $g(x) \geq \sigma(x)$. Given a *horizon* $N \in \mathbb{N}$, let us define the *cost function* and the *value function*, respectively, as

$$J_N(x, \mathbf{u}) := \sum_{k=0}^{N-1} \ell(\psi(k, x, \mathbf{u}), u_k) + g(\psi(N, x, \mathbf{u})), \;\; V_N(x) := \inf_{\mathbf{u}} J_N(x, \mathbf{u}) \quad (7)$$

where $\psi(k, x, \mathbf{u})$ is the *solution* to system (1) at time k, starting from the initial condition x, evolved under the influence of the input sequence $\mathbf{u}$. The above optimization is over the set of *admissible* input sequences, i.e. input sequences with each element residing in $\mathcal{U}$. In order to keep the discussion simple, we make the following assumption. (A less restrictive set of assumptions for a more general setting can be found in [9].)

Assumption 3.1 *For all $N \in \mathbb{N}$ and $x \in \mathbb{X}$ a minimizing input sequence $\mathbf{u}$ satisfying $V_N(x) = J_N(x, \mathbf{u})$ exists. V_N is continuous on $\mathbb{X}$ and there exists $L > 0$ such that $V_N(x) \leq L\sigma(x)$ for all $x \in \mathbb{X}$ and $N \in \mathbb{N}$.*

Given a horizon N, for $x \in \mathbb{X}$ we let the *MPC-generated feedback* $\kappa_N(x) := u_0$ where u_0 is the first element of an input sequence satisfying $V_N(x) = J_N(x, \mathbf{u})$. In this setting the following result ensues (see [9] for details).

Theorem 1. *Under Assumption 3.1 there exists $L > 0$ such that, for all horizon N, the value function V_N is continuous and satisfies*

$$\sigma(x) \leq V_N(x) \leq L\sigma(x) \qquad \forall x \in \mathbb{X} \ .$$

Moreover, for each $\rho \in (0, 1)$ there exists $n_\circ \in \mathbb{N}$ such that

$$V_N(f(x, \kappa_N(x))) - V_N(x) \leq -\rho\sigma(x) \qquad \forall x \in \mathbb{X} \ , \;\; N \geq n_\circ \ .$$

In particular, for N sufficiently large, the set $\mathcal{A}$ is asymptotically stable with basin of attraction $\mathcal{X}$.

4 Modified MPC to Decrease Sensitivity to Measurement Noise

4.1 MPC with Memory

Introduction

When using MPC for stabilization, one simple remedy to the robustness problem discussed in Section 2 seems to be to increase the so called execution horizon. That is, instead of applying the first element of an optimal input sequence and then measuring the state after one step to compute a new optimal input sequence for the new initial condition, one could apply the first $N_{\mathrm{e}} \geq 2$ elements of an optimal input sequence (in an open-loop fashion) before taking a new measurement and optimization. By doing so, if the state is close to where it is most vulnerable to measurement noise, before the next measurement it can be carried sufficiently far away (by choosing a large enough N_{e}) from that location. However, this method may be deleterious for certain applications where the conditions change quickly. This presents a trade-off between wanting to be robust to measurement noise and wanting to react quickly when conditions actually change. A compromise can be attained if one augments the state of the system with a memory variable that keeps record of previous decisions (calculations). With memory, the algorithm can be made to have preference over its previous decisions and the state can still be monitored at each step in order to take action against that preference if necessary or profitable.

Algorithm Description

To be more precise, choose the *buffer gain* $\mu > 1$ and a *memory horizon* $M \in \mathbb{N}$. Define $\Omega := \{\omega_1, \ldots, \omega_M\}$, $\omega_i \in \mathcal{U}$. Given $x \in \mathbb{X}$, let (admissible) input sequences $\mathbf{v} = \{v_0, v_1, \ldots\}$ and $\mathbf{w} = \{w_0, w_1, \ldots\}$ be defined as $\mathbf{v} := \operatorname*{argmin}_{\mathbf{u}} J_N(x, \mathbf{u})$ and

$$\mathbf{w} := \operatorname*{argmin}_{\mathbf{u}} J_N(x, \mathbf{u}) \quad \text{subject to} \quad u_{i-1} = \omega_i \quad \forall i \in \{1, \ldots, M\} .$$

Define

$$W_N(x, \Omega) := \inf_{\mathbf{u}} J_N(x, \mathbf{u}) \quad \text{subject to} \quad u_{i-1} = \omega_i \quad \forall i \in \{1, \ldots, M\}$$

and

$$\bar{\kappa}_N(x, \Omega) := \begin{cases} v_0 & \text{if} \quad W_N(x, \Omega) > \mu V_N(x) \\ w_0 & \text{if} \quad W_N(x, \Omega) \leq \mu V_N(x) \end{cases}$$

$$\bar{\pi}_N(x, \Omega) := \begin{cases} \{v_1, \ldots, v_M\} & \text{if} \quad W_N(x, \Omega) > \mu V_N(x) \\ \{w_1, \ldots, w_M\} & \text{if} \quad W_N(x, \Omega) \leq \mu V_N(x) \end{cases}$$

Note that when $W_N(x,\, \Omega) \le \mu V_N(x)$, we have $\bar{\kappa}_N(x,\, \Omega) = \omega_1$ and $\pi_N(x,\, \Omega) = \{\omega_2,\, \omega_3,\, \ldots,\, \omega_M,\, w_M\}$. The closed loop generated by this algorithm is

$$x^+ = f(x,\, \bar{\kappa}_N(x,\, \Omega)) \tag{8}$$

$$\Omega^+ = \pi_N(x,\, \Omega)\ . \tag{9}$$

We use $\psi(k,\, x,\, \Omega,\, \bar{\kappa}_N)$ to denote the solution to (8).

Theorem 2. *Let Assumption 3.1 hold. For each $\rho \in (0,\, 1)$ there exist $n_\circ \in \mathbb{N}$ and positive real numbers K and α such that for all $x \in \mathbb{X}$ and Ω*

$$W_N(f(x,\, \bar{\kappa}_N(x,\, \Omega)),\, \pi_N(x,\, \Omega)) - W_N(x,\, \Omega) \le -\rho\sigma(x) \tag{10}$$

$$\sigma(\psi(k,\, x,\, \Omega,\, \bar{\kappa}_N)) \le K\sigma(x)\exp(-\alpha k) \qquad \forall k \in \mathbb{N} \tag{11}$$

for all horizon N and memory horizon M satisfying $N \ge M + n_\circ$.

Robustness with Respect to Measurement Noise

Let us now comment on the possible extra robustness that the MPC with memory algorithm may bring to the stability of a closed loop. Suppose the stability of the closed loop obtained by standard MPC has some robustness with respect to (bounded) measurement noise characterized as (perhaps for x in some compact set)

$$V_N(f(x,\, \kappa_N(x+e))) - V_N(x) \le -\sigma(x)/2 + \alpha_v|e|$$

where N is large enough and $\alpha_v > 0$. Let us choose some $\mu > 1$. Let us be given some $\Omega = \{\omega_1,\, \ldots,\, \omega_M\}$. Then it is reasonable to expect for M and $N - M$ sufficiently large, at least for systems such as that with a disjoint attractor, that

$$W_N(f(x,\, \omega_1),\, \pi_N(x+e,\, \Omega)) - W_N(x,\, \Omega) \le -\sigma(x)/2 + \alpha_w|e|$$

with $\alpha_w > 0$ (much) smaller than α_v as long as $W_N(x,\, \Omega)$ is not way far off from $V_N(x)$, say $W_N(x,\, \Omega) \le 2\mu V_N(x)$. Now consider the closed loop (8)-(9) under measurement noise. Suppose $W_N(x+e,\, \Omega) \le \mu V_N(x+e)$. Then $\bar{\kappa}_N(x+e,\, \Omega) = \omega_1$. For μ sufficiently large it is safe to assume $W_N(x,\, \Omega) \le 2\mu V_N(x)$. Therefore we have

$$W_N(f(x,\, \bar{\kappa}_N(x+e,\, \Omega)),\, \pi_N(x+e,\, \Omega)) - W_N(x,\, \Omega) \le -\sigma(x)/2 + \alpha_w|e|\ .$$

Now consider the other case where $W_N(x+e,\, \Omega) > \mu V_N(x+e)$. Then define $\widetilde{\Omega} := \{v_0,\, \ldots,\, v_{M-1}\}$ where $\{v_0,\, v_1,\, \ldots\} =: \mathbf{v}$ and $V_N(x+e) = J_N(x+e,\, \mathbf{v})$. Note then that $W_N(x+e,\, \widetilde{\Omega}) = V_N(x+e)$ and it is safe to assume $W_N(x,\, \widetilde{\Omega}) \le 2\mu V_N(x)$ as well as $W_N(x,\, \widetilde{\Omega}) \le W_N(x,\, \Omega)$ for μ large enough. Note finally that $\bar{\kappa}_N(x+e,\, \Omega) = \bar{\kappa}_N(x+e,\, \widetilde{\Omega}) = v_0$ and $\pi_N(x+e,\, \Omega) = \pi_N(x+e,\, \widetilde{\Omega})$ in this case. Hence

$$\begin{aligned} W_N(f(x,\, \bar{\kappa}_N(x+e,\, \Omega)),\, \pi_N(x+e,\, \Omega)) - W_N(x,\, \Omega) \\ &= W_N(f(x,\, \bar{\kappa}_N(x+e,\, \widetilde{\Omega})),\, \pi_N(x+e,\, \widetilde{\Omega})) - W_N(x,\, \Omega) \\ &\le W_N(f(x,\, \bar{\kappa}_N(x+e,\, \widetilde{\Omega})),\, \pi_N(x+e,\, \widetilde{\Omega})) - W_N(x,\, \widetilde{\Omega}) \\ &\le -\sigma(x)/2 + \alpha_w|e|\ . \end{aligned}$$

The robustness of the closed loop is therefore enhanced.

4.2 MPC with Logic

Algorithm Description

The modification of the algorithm explained in the previous section aims to make the control law more *decisive*. In this section we take a different path that will have a similar effect. We augment the state with a logic (or index) variable q in order for the closed loop to adopt a hysteresis-type behavior. We begin by formally stating the procedure.

For each $q \in \{1,\, 2,\, \ldots,\, \bar{q}\} =: \mathcal{Q}$ let $\sigma_q : \mathbb{R}^n \to \mathbb{E}_{\ge 0}$ be a state measure with the following properties: (i) $\sigma_q(x) \in (0,\, \infty)$ for $x \in \mathbb{X}_q \setminus \mathcal{A}$, and $\sigma_q(x) = \infty$ for $x \in \mathbb{R}^n \setminus \mathbb{X}_q$, (ii) is continuous on $\mathbb{X}_q$, (iii) $\sigma_q(x)$ blows up either as x gets unbounded or approaches to the border of $\mathbb{X}_q$, and finally (iv) $\sigma_q(x) \ge \sigma(x)$. We then let $\ell_q : \mathbb{R}^n \times \mathcal{U} \to \mathbb{E}_{\ge 0}$ be our *q-stage cost* satisfying $\ell_q(x,\, u) \ge \sigma_q(x)$ and $g_q : \mathbb{R}^n \to \mathbb{E}_{\ge 0}$ *q-terminal cost* satisfying $g_q(x) \ge \sigma_q(x)$. We let $\bigcup_{q\in\mathcal{Q}} \mathbb{X}_q = \mathbb{X}$. Given a horizon $N \in \mathbb{N}$, we define, respectively, the *q-cost function* and the *q-value function*

$$J_N^q(x,\, \mathbf{u}) := \sum_{k=0}^{N-1} \ell_q(\psi(k,\, x,\, \mathbf{u}),\, u_k) + g(\psi(N,\, x,\, \mathbf{u})), \quad V_N^q(x) := \inf_{\mathbf{u}} J_N^q(x,\, \mathbf{u})\ .$$

We make the following assumption on V_N^q which is a slightly modified version of Assumption 3.1.

Assumption 4.1 *For all $N \in \mathbb{N}$, $q \in \mathcal{Q}$, and $x \in \mathbb{X}_q$ a minimizing input sequence $\mathbf{u}$ satisfying $V_N^q(x) = J_N^q(x,\, \mathbf{u})$ exists. V_N^q is continuous on $\mathbb{X}_q$. For each $q \in \mathcal{Q}$ there exist $L_q > 0$ such that $V_N^q(x) \le L_q\sigma_q(x)$ for all $x \in \mathbb{X}_q$ and $N \in \mathbb{N}$. There exists $L > 0$ such that for each $x \in \mathbb{X}$ there exists $q \in \mathcal{Q}$ such that $V_N^q(x) \le L\sigma(x)$ for all $N \in \mathbb{N}$.*

Let $\mu > 1$. Given $x \in \mathbb{X}$, let the input sequence $\mathbf{v}^q := \{v_0^q,\, v_1^q,\, \ldots\}$ be

$$\mathbf{v}^q := \operatorname*{argmin}_{\mathbf{u}} J_N^q(x,\, \mathbf{u})\ .$$

Let $q^* := \operatorname*{argmin}_{\varrho\in\mathcal{Q}} V_N^\varrho(x)$. Then we define

$$\tilde{\kappa}_N(x,\, q) := \begin{cases} v_0^{q^*} & \text{if} \quad V_N^q(x) > \mu V_N^{q^*}(x) \\ v_0^q & \text{if} \quad V_N^q(x) \le \mu V_N^{q^*}(x) \end{cases},\ \theta_N(x,\, q) := \begin{cases} q^* & \text{if} \quad V_N^q(x) > \mu V_N^{q^*}(x) \\ q & \text{if} \quad V_N^q(x) \le \mu V_N^{q^*}(x) \end{cases}$$

Let the closed loop generated by the algorithm be

$$x^+ = f(x,\, \tilde{\kappa}_N(x,\, q)) \tag{12}$$
$$q^+ = \theta_N(x,\, q)\ . \tag{13}$$

We use $\psi(k,\, x,\, q,\, \tilde{\kappa}_N)$ to denote the solution to (12).

Theorem 3. *Let Assumption 4.1 hold. For each $\rho \in (0,\, 1)$ there exist $n_\circ \in \mathbb{N}$ and positive real numbers K and α such that for all $x \in \mathbb{X}$ and $q \in \mathcal{Q}$*

$$V_N^{\theta_N(x,\,q)}(f(x,\, \tilde{\kappa}_N(x,\, q))) - V_N^q(x) \le -\rho\sigma(x) \tag{14}$$
$$\sigma(\psi(k,\, x,\, q,\, \tilde{\kappa}_N)) \le K\sigma(x)\exp(-\alpha k) \qquad \forall k \in \mathbb{N} \tag{15}$$

for all horizon $N \ge n_\circ$.

Robustness with Respect to Measurement Noise

We now discuss the robustness of stability of closed loops generated by MPC with logic. By Assumption 4.1, for some large enough fixed horizon N and for all $q \in \mathcal{Q}$ and $x \in \mathbb{X}_q$ it can be shown that $V_N^q(f(x,\, \tilde{\kappa}_N(x, q))) - V_N^q(x) \le -\sigma_q(x)/2$. For the analysis it makes no difference whether V_N^q is coming from an optimization problem or not. Therefore we might just as well consider the case where we have a number of control Lyapunov functions V^q active on sets $\mathbb{X}_q$ with associated feedbacks κ_q satisfying

$$V^q(f(x,\, \kappa_q(x))) - V^q(x) \le -\sigma_q(x)/2$$

for each $x \in \mathbb{X}_q$. Suppose each of the closed loops $x^+ = f(x,\, \kappa_q(x))$ has some degree of robustness characterized by (maybe for x in some compact set)

$$V^q(f(x,\, \kappa_q(x+e))) - V^q(x) \le -\sigma_q(x)/2 + \alpha_q|e|$$

where $\alpha_q > 0$. Now let us compound all these individual systems into a single one by picking $\mu > 1$ and with a switching strategy $q^+ = \theta(x,\, q)$ where θ is defined parallel to θ_N above. In the presence of measurement noise, the closed loop will be

$$x^+ = f(x,\, \kappa_{\theta(x+e,\, q)}(x))$$
$$q^+ = \theta(x + e,\, q)\ .$$

Suppose at some point x we have $\theta(x + e,\, q) = p \neq q$. That means $V^q(x + e) > \mu V^p(x + e)$. When $\mu > 1$ is large enough it is safe to assume, thanks to the continuity of the Lyapunov functions, that $V^q(x) \ge V^p(x)$ since e will be relatively small. Therefore

$$\begin{aligned} V^p(f(x,\, \kappa_p(x+e))) - V^q(x) &\le V^p(f(x,\, \kappa_p(x+e))) - V^p(x) \\ &\le -\sigma_p(x)/2 + \alpha_p|e| \\ &\le -\sigma(x)/2 + \bar{\alpha}|e| \end{aligned}$$

where $\bar{\alpha} := \max_q\{\alpha_q\}$. Therefore if we adopt $V(x, q) := V^q(x)$ as the Lyapunov function for our closed loop generated by the logic algorithm we can write

$$V(f(x, \kappa_{\theta(x+e, q)}(x+e)), \theta(x+e, q)) - V(x, q) \leq -\sigma(x)/2 + \bar{\alpha}|e|$$

for all x and q. Roughly speaking, the strength of robustness of the compound closed loop will be no less than that of the "weakest" individual system, provided that the buffer gain μ is high enough.

Figure 2 depicts the level sets of two Lyapunov functions with minima at two distinct target points. The sets $\{x : V_N^1(x)/V_N^2(x) = \mu\}$ and $\{x : V_N^2(x)/V_N^1(x) = \mu\}$ are indicated by dotted curves. The robustness margin with respect to measurement noise is related to the separation between these curves. A possible closed-loop trajectory in the absence of measurement noise is indicated by the dashed curve. Note that there is more than one switch before the trajectory gets close to one of the two target points.

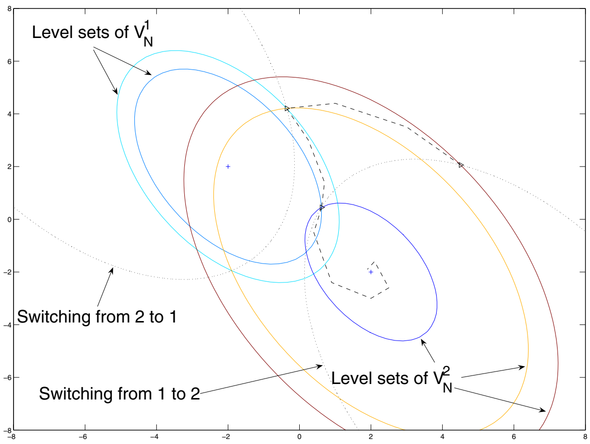

Fig. 2. Level sets of V_N^q for $q \in \{1, 2\}$. The dotted curves are the sets $\{x : V_N^1(x)/V_N^2(x) = \mu\}$ and $\{x : V_N^2(x)/V_N^1(x) = \mu\}$. The triangles represent the state at the instants when a switching occurs. The dashed line represents a piece of the solution starting at $x = (4.5, 2.1)$, the rightmost triangle.

4.3 Discussion

The two schemes offered have different advantages and disadvantages. Preference would depend on the particular application. However, MPC with memory is easier to employ in the sense that it is a minor modification to the standard algorithm. The difficulty is the determination of the design parameters M and μ; this determination is not obvious. For example, it is not true in general that the larger μ or M are, the more robustness the system has. It may be best to choose them from a range and that range possibly depends on the system and the other MPC related design parameters such as ℓ, g, and N. In the logic case, it is in general not trivial to obtain functions V^q, but it is true that a larger μ will yield more robustness to measurement error, or at least it will not degrade robustness. However, the larger μ, the longer it may take for the closed loop to converge to the desired attractor. Also, a very large μ could make the system incapable of adapting to large changes in conditions.

5 Illustrations of Modified MPC Algorithms

5.1 Pendulum Swing Up

Here we consider the problem of swinging up a pendulum and stabilizing its inverted equilibrium. The continuous-time model of the system after an input feedback transformation ($\dot{x} = F(x, v)$ where $x \in \mathbb{R}^2$, $v \in \mathbb{R}$) and normalization is $\dot{x}_1 = x_2$; $\quad \dot{x}_2 = \sin(x_1) - \cos(x_1)v$, where x_1 is the angle of the pendulum (0 at the upright position) and x_2 is the angular velocity. Following [20, Ex. 8.3], we design three different feedback laws $v_1(\cdot)$, $v_2(\cdot)$, $v_3(\cdot)$ for the system. In [20], each of these control laws are activated in a different prespecified region of the state space to perform the swing-up (the design purpose of $v_1(\cdot)$, $v_2(\cdot)$, $v_3(\cdot)$ is to kick the system from the resting condition, to pump energy into the system, and to stabilize the inverted pendulum to the upright position, respectively). Given a sampling period $T > 0$, for each $u \in \{1, 2, 3\}$ let $x^+ = f(x, u)$ be the discrete-time model of the closed loop $\dot{x} = F(x, v_u(x))$ obtained via integration over an interval of length T, i.e. $f(x, u) = \phi(T)$ where $\phi(\cdot)$ is the solution of $\dot{x} = F(x, v_u(x))$ starting at $\phi(0) = x$. We can now use MPC to decide the swing up strategy.

We construct the stage cost for standard MPC by adding the kinetic and potential energy of the pendulum. We also include a term in the stage cost that penalizes the control law during the continuous-time horizon to avoid large control efforts. The cost function is periodic in x_1 with period 2π and therefore, there exists a surface on the state space $x_1 - x_2$ where on one side the algorithm tries to reach the upright position rotating the pendulum clockwise and on the other side rotating the pendulum counterclockwise. For one such particular cost, the surface and two different trajectories in opposite directions starting close to the surface are given in Fig. 3(a). As discussed in Section 2, the closed-loop system is vulnerable to small measurement noise in the vicinity of that surface when T is small.

The vulnerability to measurement noise mentioned above can be resolved via the approach discussed in Section 4.2. Despite the fact that $x_1 = 2\pi k$, $k \in \{0, \pm 1, \pm 2, \ldots\}$, correspond to the same physical location, one can construct two stage costs, namely ℓ_q for $q \in \{1, 2\}$, that are not periodic in x_1 such that ℓ_1 vanishes at $x = (0, 0, 0, 0)$ and positive elsewhere and ℓ_2 vanishes at $x = (2\pi, 0, 0, 0)$ and positive elsewhere. By doing so we can attain a robustness margin that does not depend on the size of sampling period T but on μ only, which can be increased to enhance robustness. Fig. 3(b) shows the switching lines for several values of μ for both possible switches ($q = 1 \to 2$, $q = 2 \to 1$). For a particular value of μ, the robustness margin is related to the separation of the

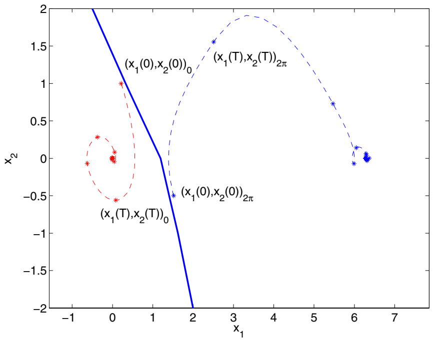

(a) Vulnerability of standard MPC to measurement noise: trajectories, denoted by $*$, starting close to the thick line approach different equilibrium points.

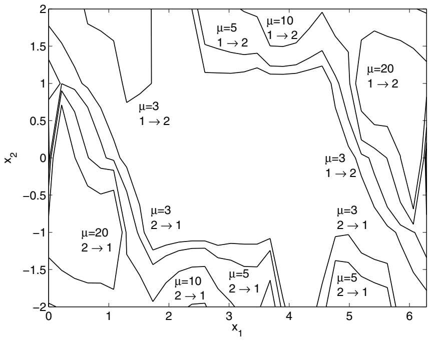

(b) MPC with logic: switching lines from $q = 1 \to 2$ and $q = 2 \to 1$ for various μ.

Fig. 3. Swing-up with standard MPC and MPC with memory

lines. The margin is independent of the sampling time T as long as NT remains constant, N being the horizon for MPC with logic. The design of an MPC with memory controller and the extension to the case of swinging up the pendulum on a cart follows directly, but due to space limitations we do not include them here.

5.2 Obstacle Avoidance with Constant Horizontal Velocity

Consider a vehicle moving on the plane $x^+ = x + \delta,\ y^+ = y + u\delta$ where $\delta > 0$ and $u \in \{-1, 1\}$ (note that this system can be thought of as sampling the system $\dot{x} = 1,\ \dot{y} = u$). Suppose that the goal for the vehicle is to avoid hitting an obstacle defined by a block of unit height centered about the horizontal axis at $x = 0$ (i.e. the vehicle must leave the region $y \in [-0.5, 0.5]$ before $x = 0$). We design a controller using MPC with logic. Let $q \in \{1, 2\}$, $\ell_1([x, y]^T, u) = \ell_2([x, -y]^T, u) = \exp(y)$, and $g(\cdot) = 0$. Since the costs are invariant on x and symmetric about the x axis, the decision lines defined by μ turn out to be horizontal lines. Let the spacing between these lines be $s(\mu)$. In this case, $s(\mu) = \ln(\mu)$, since $V_N^1([x, y]^T) = \mu V_N^2([x, y]^T)$ when $y = \frac{\ln(\mu)}{2}$ and $V_N^2([x, y]^T) = \mu V_N^1([x, y]^T)$ when $y = -\frac{\ln(\mu)}{2}$ for any N.

Note that when $\mu = 1$ (or $s(\mu) = 0$) MPC with logic is equivalent to the standard MPC algorithm implemented using the stage cost $\ell([x, y]^T, u) = \min\{\ell_1([x, y]^T, u), \ell_2([x, y]^T, u)\}$. As μ is increased, the spacing $s(\mu)$ increases. Table 1 shows the average number of switches and the total number of crashes for $50,000$ runs of the system. The initial conditions are set to be $x(0) = -1.5$ and $y(0)$ normally distributed (though kept within $(-1, 1)$) around $y = 0$. The noise is uniformly distributed in $[-0.8, 0.8]$. The key variables of comparison are the spacing of the decision lines $s(\mu)$ and the sampling time δ. With the increased

Table 1. Simulations of system with differing decision line spacing and sampling time for uniformly distributed measurement noise $e \in [-0.8, 0.8]$. Each datum is generated by 50,000 runs starting at $x(0) = -1.5$ and $y(0)$ normally distributed constrained to $(-1, 1)$. "TC" is total number of crashes and "AS" is average number of switches.

δ	0.1		0.06		0.03		0.01		0.006		0.003		0.001	
$s(\mu)$	TC	AS	TC	AS	TC	AS	TC	AS	TC	AS	TC	AS	TC	AS
0.00	5110	2.61	5444	4.36	5791	8.73	5927	26.1	5922	43.8	6107	88.3	6125	263
0.25	3248	1.75	3716	2.88	3862	5.72	4197	17.2	4334	28.9	4375	57.3	4144	169
0.50	1575	1.13	1875	1.76	2253	3.38	2549	10.0	2529	16.7	2625	33.2	2638	101
0.75	428	0.70	609	1.01	863	1.80	1102	5.00	1140	8.26	1155	16.4	1189	48.6
1.00	47	0.46	51	0.57	110	0.86	241	2.05	274	3.21	276	6.14	298	17.7
1.25	1	0.34	2	0.39	1	0.48	6	0.73	9	0.94	8	1.55	17	3.92
1.50	0	0.25	0	0.29	0	0.35	0	0.42	0	0.45	0	0.48	0	0.54

spacing for a given sampling time, there are fewer crashes, as expected, and the trajectories contain fewer switches. The number of switches can be thought of as a measure of the sensitivity to measurement noise. As the sampling time is decreased, the system also becomes more sensitive to measurement noise due to the smaller movements of the system making it difficult to escape the neighborhood of the horizontal axis.

For this system, a crash-free bound on the measurement noise (that solely depends on μ) can be calculated as follows.

Claim. Suppose the MPC with logic controller is implemented with the cost functions ℓ_1, ℓ_2. If the buffer gain $\mu > 1$, the measurement noise is bounded by $\frac{s(\mu)}{2}$, and the horizontal component of the state $x < -\left(\frac{1+s(\mu)}{2}\right)$ then the system will not crash due to measurement noise.

Note that the bound in Claim 5.2 does not depend on the sampling time δ. Hence, the given controller yields a robustness margin independent of δ. For this system, increasing the buffer gain will always increase the robustness margin. However, this may not work on other systems. Increasing the buffer gain too much can cause a system to become obstinate rather than decisive. Choosing the buffer gain then will be very dependent on the task that the system is required to perform. A balance must be made between ignoring (usually small) measurment error and responding to (relatively large) changes in task conditions.

5.3 Avoiding Moving Obstacles

Let the dynamics of a vehicle and obstacle be $x^+ = x + u$ and $z^+ = z + v$, respectively, where $x \in \mathbb{R}^2$, $u \in \{-1, 0, 1\} \times \{-1, 0, 1\}$, and $v \in \{-1, 0, 1\}$. We fix the vertical displacement of the obstacle $h > 0$, and constrain the horizontal displacement to $z \in [-1, 1]$. The goal of the vehicle is to reach some target while avoiding the obstacle whose goal is to reach the vehicle. Both of the agents are considered as single points in $\mathbb{R}^2$ and run MPC to achieve their goals as follows. The stage cost of the vehicle puts a high penalty on the current location of the obstacle and gradually vanishes at the target. The stage cost of the obstacle vanishes whenever $x_1 = z$ and is positive elsewhere.

Applying standard MPC with the vehicle and the obstacle initially aligned vertically at zero horizontal position, the obstacle is able to prevent the vehicle from converging to its target. Suppose the vehicle decides to move in the increasing x_1 direction to avoid the obstacle from the right. The obstacle will follow the vehicle with one step of delay. At some point, it will become necessary for the vehicle to change its course since the optimal path, now that the obstacle has moved, is now to the the left of the obstacle. Hence the vehicle can get stuck possibly as shown in Fig. 4(a).

With the MPC with memory approach described in Section 4.1, the problem can be resolved. Using the same stage cost, $M = 5$, and $\mu = 1.4$, the vehicle avoids the obstacle. The sequence in memory is effectively used when the obstacle is at $(2, 3)$: the vehicle stays within his initial course of passing the obstacle from the right as shown in Fig. 4(b). Similar results were obtained with MPC with logic using two symmetric respect to x_1 stage cost functions that allow the vehicle avoid the obstacle from the left and from the right, respectively, but the results are omitted because of space limitations.

Note that the moving obstacle can be thought of as noise for the measurement of the vehicle's distance to a static obstacle. Since the displacement of the obstacle has magnitude equal to one, the "measurement noise" for the vehicle is rather large.

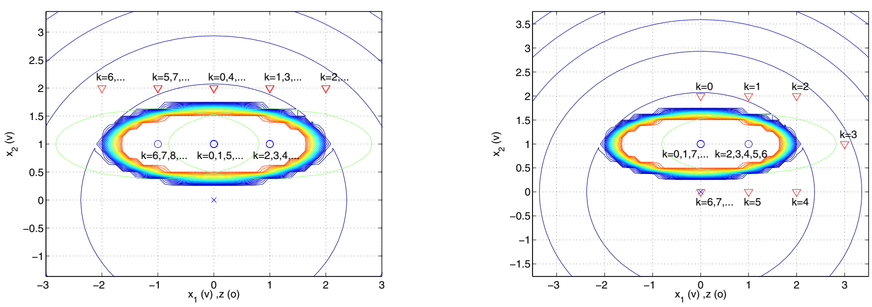

(a) Standard MPC: target acquisition failed. Vehicle's vertical position is always equal to two. The thick ellipse corresponds to the bump created by the obstacle in the vehicle's stage cost.

(b) MPC with memory: target acquisition successful. Vehicle goes around the obstacle to the right. The ellipse centered at $(1, 1)$ denotes the bump in the vehicle's stage cost when the obstacle moves.

Fig. 4. Comparison of standard MPC and MPC with memory using identical stage cost and $N = 6$. The vehicle is represented by ∇, the obstacle by $\circ$, and the target by $\times$ (the origin). The discrete time is denoted by k and the level sets of the stage cost for the vehicle are plotted for the initial condition.

References

[1] A. Bemporad, F. Borrelli, and M. Morari. Optimal controllers for hybrid systems: Stability and piecewise linear explicit form. *Proc. IEEE Conf. on Decision and Control*, Vol. 2, pp. 1810–1815, (2000).

[2] F.H. Clarke, S. Yu, S. Ledyaev, L. Rifford, and R.J. Stern. Feedback stabilization and Lyapunov functions. *SIAM Journal on Control and Optimization*, Vol. 39, pp. 25–48, (2000).

[3] P. Collins. A trajectory-space approach to hybrid systems. *Proc. of the International Symposium on Mathematical Theory of Networks and Systems*, (2004).

[4] J-M. Coron and L. Rosier. A relation between continuous time-varying and discontinuous feedback stabilization.*Journal of Mathematical Systems, Estimation, and Control*, Vol. 4, pp. 67–84, (1994).

[5] R. Goebel, J.P. Hespanha, A.R. Teel, C. Cai, and R.G. Sanfelice. Hybrid systems: Generalized solutions and robust stability.*Proc. NOLCOS*, Vol. 1, pp. 1–12, (2004).

[6] R. Goebel and A.R. Teel. Results on solution sets to hybrid systems with applications to stability theory. *Proc. American Control Conf.*, Vol. 1, pp. 557–562 (2005).

[7] R. Goebel and A.R. Teel. Solutions to hybrid inclusions via set and graphical convergence. *Automatica*, Available online 20 February, (2006).

[8] G. Grimm, M.J. Messina, S.E. Tuna, and A.R. Teel. Examples when nonlinear model predictive control is nonrobust. *Automatica*, Vol. 40, pp. 1729–1738, (2004).

[9] G. Grimm, M.J. Messina, S.E. Tuna, and A.R. Teel. Model predictive control: for want of a local control Lyapunov function, all is not lost. *IEEE Trans. Auto. Control*, Vol. 50, pp. 546–558, (2005).

[10] W.P.M.H Heemels, B. De Schutter, and A. Bemporad. Equivalence of hybrid dynamical models. *Automatica*, Vol. 37, pp.1085–1091, (2001).

[11] Y. Kuwata and J. How. Receding horizon implementation of MILP for vehicle guidance. In *Proc. American Control Conf.*, pp. 2684–2685, (2005).

[12] T. Lapp and L. Singh. Model predictive control based trajectory optimization for nap-of-the-earth (NOE) flight including obstacle avoidance. *Proc. American Control Conf.*, Vol. 1, pp. 891–896, (2004).

[13] J. Lygeros, K.H. Johansson, S.N. Simić, J. Zhang, and S.S. Sastry. Dynamical properties of hybrid automata. *IEEE Trans. Auto. Control*, Vol. 48, pp. 2–17, (2003).

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

[15] M.J. Messina, S.E. Tuna, and A.R. Teel. Discrete-time certainty equivalence output feedback, pp. allowing discontinuous control laws including those from model predictive control. *Automatica*, Vol. 41, pp. 617–628, (2005).

[16] P. Mhaskar, N.H. El-Farrah, and P.D. Christofides. Robust hybrid predictive control of nonlinear systems. *Automatica*, Vol. 41, pp. 209–217, (2005).

[17] H. Michalska and D.Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Trans. Auto. Control*, Vol. 38, pp. 1623–1633, (1993).

[18] P. Ogren and N.E. Leonard. A convergent dynamic window approach to obstacle avoidance. *IEEE Trans. on Robotics*, Vol. 21, pp. 188–195, (2005).

[19] C. Prieur. Perturbed hybrid systems, applications in control theory. *Nonlinear and Adaptive Control NCN4 2001 (Lecture Notes in Control and Information Sciences Vol. 281)*, pp. 285–294, (2003).

[20] R. G. Sanfelice, R. Goebel, and A.R. Teel. Invariance principles for hybrid systems with connections to detectability and asymptotic stability. Submitted to IEEE Trans. Auto. Control, (2005).

[21] E.D. Sontag. Clocks and insensitivity to small measurement errors. *Control, Optimisation and Calculus of Variations*, Vol. 4, pp. 537–557, (1999).

[22] A. van der Schaft and H. Schumacher. *An Introduction to Hybrid Dynamical Systems (Lecture Notes in Control and Information Sciences Vol. 251)*, (2000).