We are IntechOpen,
the world's leading publisher of
Open Access books
Built by scientists, for scientists

6,900
Open access books available

186,000
International authors and editors

200M
Downloads

154
Countries delivered to

Our authors are among the

TOP 1%
most cited scientists

12.2%
Contributors from top 500 universities

BOOK CITATION INDEX
CLARIVATE ANALYTICS
INDEXED

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Robust Adaptive Model Predictive Control of Nonlinear Systems

Darryl DeHaan and Martin Guay
Dept. Chemical Engineering, Queen's University
Canada

1. Introduction

When faced with making a decision, it is only natural that one would aim to select the course of action which results in the "best" possible outcome. However, the ability to arrive at a decision necessarily depends upon two things: a well-defined notion of what qualities make an outcome desirable, and a *previous* decision[1] defining to what extent it is necessary to characterize the quality of individual candidates before making a selection (i.e., a notion of when a decision is "good enough"). Whereas the first property is required for the problem to be well defined, the later is necessary for it to be tractable.

The process of searching for the "best" outcome has been mathematically formalized in the framework of optimization. The typical approach is to define a scalar-valued *cost function*, that accepts a decision candidate as its argument, and returns a quantified measure of its quality. The decision-making process then reduces to selecting a candidate with the lowest (or highest) such measure.

1.1 The Emergence of Optimal Control

The field of "control" addresses the question of how to manipulate an *input u* in order to drive the *state x* of a dynamical system

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

to some desired target. Ultimately this task can be viewed as decision-making, so it is not surprising that it lends itself towards an optimization-based characterization. Assuming that one can provide the necessary metric for assessing quality of the trajectories generated by (1), there exists a rich body of "optimal control" theory to guide this process of decision-making. Much of this theory came about in the 1950's and 60's, with Pontryagin's introduction of the Minimum (a.k.a. Maximum) Principle Pontryagin (1961), and Bellman's development of Dynamic Programming Bellman (1952; 1957). (This development also coincided with landmark results for linear systems, pioneered by Kalman Kalman (1960; 1963), that are closely related). However, the roots of both approaches actually extend back to the mid-1600's, with the inception of the calculus of variations.

---

[1] The recursiveness of this definition is of course ill-posed until one accepts that at some level, every decision is ultimately predicated upon underlying assumptions, accepted entirely in faith.

The tools of optimal control theory provide useful benchmarks for characterizing the notion of "best" decision-making, as it applies to control. However applied directly, the tractability of this decision-making is problematic. For example, Dynamic Programming involves the construction of a $n-$dimensional surface that satisfies a challenging nonlinear partial differential equation, which is inherently plagued by the so-called *curse of dimensionality*. This methodology, although elegant, remains generally intractable for problems beyond modest size. In contrast, the Minimum Principle has been relatively successful for use in off-line trajectory planning, when the initial condition of (1) is known. Although it was suggested as early as 1967 in Lee & Markus (1967) that a stabilizing feedback $u = k(x)$ could be constructed by continuously re-solving the calculations online, a tractable means of doing this was not immediately forthcoming.

1.2 Model Predictive Control as Receding-Horizon Optimization

Early development (Richalet et al. (1976),Richalet et al. (1978),Cutler & Ramaker (1980)) of the control approach known today as Model Predictive Control (MPC) originated in the process control community, and was driven much more by industrial application than by theoretical understanding. Modern theoretical understanding of MPC, much of which developed throughout the 1990's, has clarified its very natural ties to existing optimal control theory. Key steps towards this development included such results as Chen & Allgöwer (1998a;b); De Nicolao et al. (1996); Jadbabaie et al. (2001); Keerthi & Gilbert (1988); Mayne & Michalska (1990); Michalska & Mayne (1993); Primbs et al. (2000), with an excellent unifying survey in Mayne et al. (2000).

At its core, MPC is simply a framework for implementing existing tools of optimal control. Taking the current value $x(t)$ as the initial condition for (1), the Minimum Principle is used as the primary basis for identifying the "best" candidate trajectory by predicting the future behaviour of the system using model (1). However, the actual quality measure of interest in the decision-making is generally the total future accumulation (i.e., over an infinite future) of a given instantaneous metric, a quantity rarely computable in a satisfactorily short time. As such, MPC only generates predictions for (1) over a finite time-horizon, and approximates the remaining infinite tail of the cost accumulation using a penalty surface derived from either a *local* solution of the Dynamic Programming surface, or an appropriate approximation of that surface. As such, the key benefit of MPC over other optimal control methods is simply that its finite horizon allows for a convenient trade-off between the online computational burden of solving the Minimum Principle, and the offline burden of generating the penalty surface.

In contrast to other approaches for constructive nonlinear controller design, optimal control frameworks facilitate the inclusion of constraints, by imposing feasibility of the candidates as a condition in the decision-making process. While these approaches can be numerically burdensome, optimal control (and by extension, MPC) provides the only real framework for addressing the control of systems in the presence of constraints - in particular those involving the state x. In practice, the predictive aspect of MPC is unparalleled in its ability to account for the risk of future constraint violation during the current control decision.

1.3 Current Limitations in Model Predictive Control

While the underlying theoretical basis for model predictive control is approaching a state of relative maturity, application of this approach to date has been predominantly limited to "slow" industrial processes that allow adequate time to complete the controller calculations. There is great incentive to extend this approach to applications in many other sectors, moti-

vated in large part by its constraint-handling abilities. Future applications of significant interest include many in the aerospace or automotive sectors, in particular constraint-dominated problems such as obstacle avoidance. At present, the significant computational burden of MPC remains the most critical limitation towards its application in these areas.

The second key weakness of the model predictive approach remains its susceptibility to uncertainties in the model (1). While a fairly well-developed body of theory has been developed within the framework of robust-MPC, reaching an acceptable balance between computational complexity and conservativeness of the control remains a serious problem. In the more general control literature, adaptive control has evolved as an alternative to a robust-control paradigm. However, the incorporation of adaptive techniques into the MPC framework has remained a relatively open problem.

2. Notational and Mathematical Preliminaries

Throughout the remainder of this dissertation, the following is assumed by default (where $s \in \mathbb{R}^s$ and S represent arbitrary vectors and sets, respectively):

- all vector norms are Euclidean, defining balls $B(s, \delta) \triangleq \{s' \mid \|s - s'\| \leq \delta\}$, $\delta \geq 0$.

- norms of matrices $\mathcal{S} \in \mathbb{R}^{m \times s}$ are assumed induced as $\|\mathcal{S}\| \triangleq \max_{\|s\|=1} \|\mathcal{S}s\|$.

- the notation $s_{[a,b]}$ denotes the entire continuous-time trajectory $s(\tau)$, $\tau \in [a, b]$, and likewise $\dot{s}_{[a,b]}$ the trajectory of its forward derivative $\dot{s}(\tau)$.

- For any set $\mathsf{S} \subseteq \mathbb{R}^s$, define

  i) its closure $\mathrm{cl}\{\mathsf{S}\}$, interior $\mathring{\mathsf{S}}$, and boundary $\partial \mathsf{S} = \mathrm{cl}\{\mathsf{S}\} \setminus \mathring{\mathsf{S}}$

  ii) its orthogonal distance norm $\|s\|_{\mathsf{S}} \triangleq \inf_{s' \in \mathsf{S}} \|s - s'\|$

  iii) a closed δ-neighbourhood $B(\mathsf{S}, \delta) \triangleq \{s \in \mathbb{R}^s \mid \|s\|_{\mathsf{S}} \leq \delta\}$

  iv) an interior approximation $\overleftarrow{B}(\mathsf{S}, \delta) \triangleq \{s \in \mathsf{S} \mid \inf_{s' \in \partial \mathsf{S}} \|s - s'\| \geq \delta\}$

  v) a (finite, closed, open) cover of S as any (finite) collection $\{\mathsf{S}^i\}$ of (open, closed) sets $\mathsf{S}^i \subseteq \mathbb{R}^s$ such that $\mathsf{S} \subseteq \cup_i \mathsf{S}^i$.

  vi) the maximal closed subcover $\overline{\mathrm{cov}}\{\mathsf{S}\}$ as the infinite collection $\{\mathsf{S}^i\}$ containing *all possible* closed subsets $\mathsf{S}^i \subseteq \mathsf{S}$; i.e., $\overline{\mathrm{cov}}\{\mathsf{S}\}$ is a maximal "set of subsets".

Furthermore, for any arbitrary function $\alpha : \mathsf{S} \to \mathbb{R}$ we assume the following definitions:

- $\alpha(\cdot)$ is C^{m+} if it is at least m-times differentiable, with all derivatives of order m yielding locally Lipschitz functions.

- A function $\alpha : \mathsf{S} \to (-\infty, \infty]$ is *lower semi-continuous* ($\mathcal{LS}$-continuous) at s if it satisfies (see Clarke et al. (1998)):
$$\liminf_{s' \to s} \alpha(s') \geq \alpha(s) \tag{2}$$

- a continuous function $\alpha : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ belongs to *class* $\mathcal{K}$ if $\alpha(0) = 0$ and $\alpha(\cdot)$ is strictly increasing on $\mathbb{R}_{>0}$. It belongs to *class* $\mathcal{K}_{\infty}$ if it is furthermore radially unbounded.

- a continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ belongs to *class* $\mathcal{KL}$ if i) for every fixed value of τ, it satisfies $\beta(\cdot, \tau) \in \mathcal{K}$, and ii) for each fixed value of s, then $\beta(s, \cdot)$ is strictly decreasing and satisfies $\lim_{\tau \to \infty} \beta(s, \tau) = 0$.

- the scalar operator $\mathrm{sat}_a^b(\cdot)$ denotes saturation of its arguments onto the interval $[a, b]$, $a < b$. For vector- or matrix-valued arguments, the saturation is presumed by default to be evaluated element-wise.

3. Brief Review of Optimal Control

The underlying assumption of optimal control is that at any time, the pointwise cost of x and u being away from their desired targets is quantified by a known, physically-meaningful function $L(x, u)$. Loosely, the goal is to then reach some target in a manner that accumulates the least cost. It is not generally necessary for the "target" to be explicitly described, since its knowledge is built into the function $L(x, u)$ (i.e it is assumed that convergence of x to any invariant subset of $\{x \mid \exists u \quad s.t. \ L(x, u) = 0\}$ is as acceptable). The following result, while superficially simple in appearance, is in fact the key foundation underlying the optimal control results of this section, and by extension all of model predictive control as well. Proof can be found in many references, such as Sage & White (1977).

Definition 3.1 (Principle of Optimality:). *If $u^*_{[t_1,t_2]}$ is an optimal trajectory for the interval $t \in [t_1, t_2]$, with corresponding solution $x^*_{[t_1,t_2]}$ to (1), then for any $\tau \in (t_1, t_2)$ the sub-arc $u^*_{[\tau,t_2]}$ is necessarily optimal for the interval $t \in [\tau, t_2]$ if (1) starts from $x^*(\tau)$.*

4. Variational Approach: Euler, Lagrange and Pontryagin

Pontryagin's Minimum principle (also known as the Maximum principle, Pontryagin (1961)) represented a landmark extension of classical ideas of variational calculus to the problem of control. Technically, the Minimum Principle is an application of the classical Euler-Lagrange and Weierstrass conditions[2] Hestenes (1966), which provide first-order *necessary* conditions to characterize extremal time-trajectories of a cost functional.[3]. The Minimum Principle therefore characterizes minimizing trajectories $(x_{[0,T]}, u_{[0,T]})$ corresponding to a constrained finite-horizon problem of the form

$$V_T(x_0, u_{[0,T]}) = \int_0^T L(x, u)\, d\tau + W(x(T)) \tag{3a}$$

$$s.t. \ \forall \tau \in [0, T]:$$

$$\dot{x} = f(x, u), \qquad x(0) = x_0 \tag{3b}$$

$$g(x(\tau)) \leq 0, \qquad h(x(\tau), u(\tau)) \leq 0, \qquad w(x(T)) \leq 0 \tag{3c}$$

where the vectorfield $f(\cdot, \cdot)$ and constraint functions $g(\cdot)$, $h(\cdot, \cdot)$, and $w(\cdot)$ are assumed sufficiently differentiable.

Assume that $g(x_0) < 0$, and, for a given $(x_0, u_{[0,T]})$, let the interval $[0, T)$ be partitioned into (maximal) subintervals as $\tau \in \cup_{i=1}^p [t_i, t_{i+1})$, $t_0 = 0$, $t_{p+1} = T$, where the interior t_i represent intersections $g < 0 \Leftrightarrow g = 0$ (i.e., the $\{t_i\}$ represent changes in the active set of g). Assuming that $g(x)$ has constant relative degree r over some appropriate neighbourhood, define the following vector of (Lie) derivatives: $N(x) \triangleq [g(x), g^{(1)}(x), \ldots g^{(r-1)}(x)]^T$, which characterizes additional tangency constraints $N(x(t_i)) = 0$ at the corners $\{t_i\}$. Rewriting (3) in multiplier form

$$V_T = \int_0^T \mathcal{H}(x, u) - \lambda^T \dot{x}\, d\tau + W(x(T)) + \mu_w w(x(T)) + \sum_i \mu_N^T(t_i) N(x(t_i)) \tag{4a}$$

$$\mathcal{H} \triangleq L(x, u) + \lambda^T f(x, u) + \mu_h h(x, u) + \mu_g g^{(r)}(x, u) \tag{4b}$$

---

[2] phrased as a fixed initial point, free endpoint problem

[3] i.e., generalizing the NLP necessary condition $\frac{\partial p}{\partial x} = 0$ for the extrema of a function $p(x)$.

overa Taking the first variation of the right-hand sides of (4a,b) with respect to perturbations in $x_{[0,T]}$ and $u_{[0,T]}$ yields the following set of conditions (adapted from statements in Bertsekas (1995); Bryson & Ho (1969); Hestenes (1966)) which necessarily must hold for V_T to be minimized:

Proposition 4.1 (Minimum Principle). *Suppose that the pair $(u^*_{[0,T]}, x^*_{[0,T]})$ is a minimizing solution of (3). Then for all $\tau \in [0,T]$, there exists multipliers $\lambda(\tau) \geq 0$, $\mu_h(\tau) \geq 0$, $\mu_g(\tau) \geq 0$, and constants $\mu_w \geq 0$, $\mu_N^i \geq 0$, $i \in \mathcal{I}$, such that*

i) *Over each interval $\tau \in [t_i, t_{i+1}]$, the multipliers $\mu_h(\tau)$, $\mu_g(\tau)$ are piecewise continuous, $\mu_N(\tau)$ is constant, $\lambda(\tau)$ is continuous, and with $(u^*_{[t_i, t_{i+1}]}, x^*_{[t_i, t_{i+1}]})$ satisfies*

$$\dot{x}^* = f(x^*, u^*), \qquad x^*(0) = x_0 \tag{5a}$$

$$\dot{\lambda}^T = \nabla_x \mathcal{H} \quad a.e., \ with \quad \lambda^T(T) = \nabla_x W(x^*(T)) + \mu_w \nabla_x w(x^*(T)) \tag{5b}$$

where the solution $\lambda_{[0,T]}$ is discontinuous at $\tau \in \{t_i\}$, $i \in \{1,3,5...p\}$, satisfying

$$\lambda^T(t_i^-) = \lambda^T(t_i^+) + \mu_N^T(t_i^+) \nabla_x N(x(t_i)) \tag{5c}$$

ii) *$\mathcal{H}(x^*, u^*, \lambda, \mu_h, \mu_g)$ is constant over intervals $\tau \in [t_i, t_{i+1}]$, and for all $\tau \in [0,T]$ it satisfies (where $\mathcal{U}(x) \triangleq \{u \mid h(x,u) \leq 0 \ and \ \left(g^{(r)}(x,u) \leq 0 \ if \ g(x) = 0\right)\}$):*

$$\mathcal{H}(x^*, u^*, \lambda, \mu_h, \mu_g) \leq \min_{u \in \mathbf{U}(x)} \mathcal{H}(x^*, u, \lambda, \mu_h, \mu_g) \tag{5d}$$

$$\nabla_u \mathcal{H}(x^*(\tau), u^*(\tau), \lambda(\tau), \mu_h(\tau), \mu_g(\tau)) = 0 \tag{5e}$$

iii) *For all $\tau \in [0,T]$, the following constraint conditions hold*

$$g(x^*) \leq 0 \qquad\qquad h(x^*, u^*) \leq 0 \qquad\qquad w(x^*(T)) \leq 0 \tag{5f}$$

$$\mu_g(\tau)g^{(r)}(x^*, u^*) = 0 \qquad \mu_h(\tau)h(x^*, u^*) = 0 \qquad \mu_w w(x^*(T)) = 0 \tag{5g}$$

$$\mu_N^T(\tau)N(x^*) = 0 \qquad \left(and \ N(x^*) = 0, \ \forall \tau \in [t_i, t_{i+1}], \ i \in \{1,3,5...p\}\right) \tag{5h}$$

The multiplier $\lambda(t)$ is called the *co-state*, and it requires solving a two-point boundary-value problem for (5a) and (5b). One of the most challenging aspects to locating (and confirming) a minimizing solution to (5) lies in dealing with (5c) and (5h), since the number and times of constraint intersections are not known a-priori.

5. Dynamic Programming: Hamilton, Jacobi, and Bellman

The Minimum Principle is fundamentally based upon establishing the optimality of a *particular* input trajectory $u_{[0,T]}$. While the applicability to offline, open-loop trajectory planning is clear, the inherent assumption that x_0 is known can be limiting if one's goal is to develop a *feedback* policy $u = k(x)$. Development of such a policy requires the consideration of *all possible* initial conditions, which results in an optimal cost *surface* $J^* : \mathbb{R}^n \to \mathbb{R}$, with an associated control policy $k : \mathbb{R}^n \to \mathbb{R}^m$. A constructive approach for calculating such a surface, referred to as *Dynamic Programming*, was developed by Bellman Bellman (1957). Just as the

Minimum Principle was extended out of the classical trajectory-based Euler-Lagrange equations, Dynamic Programming is an extension of classical Hamilton-Jacobi field theory from the calculus of variations.

For simplicity, our discussion here will be restricted to the unconstrained problem:

$$V^*(x_0) = \min_{u_{[0,\infty)}} \int_0^\infty L(x,u)\, d\tau \tag{6a}$$

$$s.t. \quad \dot{x} = f(x,u), \qquad x(0) = x_0 \tag{6b}$$

with locally Lipschitz dynamics $f(\cdot, \cdot)$. From the Principle of Optimality, it can be seen that (6) lends itself to the following recursive definition:

$$V^*(x(t)) = \min_{u[t,t+\Delta t]} \left\{ \int_t^{t+\Delta t} L(x(\tau), u(\tau)) d\tau + V^*(x(t+\Delta t)) \right\} \tag{7}$$

Assuming that V^* is differentiable, replacing $V^*(x(t+\Delta t))$ with a first-order Taylor-series and the integrand with a Riemannian sum, the limit $\Delta t \to 0$ yields

$$0 = \min_u \left\{ L(x,u) + \frac{\partial V^*}{\partial x} f(x,u) \right\} \tag{8}$$

Equation (8) is one particular form of what is known as the Hamilton-Jacobi-Bellman (HJB) equation. In some cases (such as $L(x,u)$ quadratic in u, and $f(x,u)$ affine in u), (8) can be simplified to a more standard-looking PDE by evaluating the indicated minimization in closed-form[4]. Assuming that a (differentiable) surface $V^* : \mathbb{R}^n \to \mathbb{R}$ is found (generally by off-line numerical solution) which satisfies (8), a stabilizing feedback $u = k_{DP}(x)$ can be constructed from the information contained in the surface V^* by simply defining[5] $k_{DP}(x) \triangleq \{u \mid \frac{\partial V^*}{\partial x} f(x,u) = -L(x,u)\}$.

Unfortunately, incorporation of either input or state constraints generally violates the assumed smoothness of $V^*(x)$. While this could be handled by interpreting (8) in the context of *viscosity solutions* (see Clarke et al. (1998) for definition), for the purposes of application to model predictive control it is more typical to simply restrict the domain of $V^* : \Omega \to \mathbb{R}$ such that $\Omega \subset \mathbb{R}^n$ is feasible with respect to the constraints.

6. Inverse-Optimal Control Lyapunov Functions

While knowledge of a surface $V^*(x)$ satisfying (8) is clearly ideal, in practice analytical solutions are only available for extremely restrictive classes of systems, and almost never for systems involving state or input constraints. Similarly, numerical solution of (8) suffers the so-called "curse of dimensionality" (as named by Bellman) which limits its applicability to systems of restrictively small size.

An alternative design framework, originating in Sontag (1983), is based on the following:

Definition 6.1. *A* control Lyapunov function (CLF) *for* (1) *is any* C^1, *proper, positive definite function* $V : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ *such that, for all* $x \neq 0$:

$$\inf_u \frac{\partial V}{\partial x} f(x,u) < 0 \tag{9}$$

---

[4] In fact, for linear dynamics and quadratic cost, (8) reduces down to the linear Ricatti equation.

[5] $k_{DP}(\cdot)$ is interpreted to incorporate a deterministic selection in the event of multiple solutions. The existence of such a u is implied by the assumed solvability of (8)

Design techniques for deriving a feedback $u = k(x)$ from knowledge of $V(\cdot)$ include the well-known "Sontag's Controller" of Sontag (1989), which led to the development of "Pointwise Min-Norm" control of the form Freeman & Kokotović (1996a;b); Sepulchre et al. (1997):

$$\min_u \gamma(u) \quad s.t. \quad \frac{\partial V}{\partial x} f(x, u) < -\sigma(x) \tag{10}$$

where γ, σ are positive definite, and γ is radially unbounded. As discussed in Freeman & Kokotović (1996b); Sepulchre et al. (1997), relation (9) implies that *there exists* a function $L(x, u)$, derived from γ and σ, for which $V(\cdot)$ satisfies (8). Furthermore, if $V(x) \equiv V^*(x)$, then appropriate selection of γ, σ (in particular that of Sontag's controller Sontag (1989)) results in the feedback $u = k_{clf}(x)$ generated by (9) satisfying $k_{clf}(\cdot) \equiv k_{DP}(\cdot)$. Hence this technique is commonly referred to as "inverse-optimal" control design, and can be viewed as a method for approximating the optimal control problem (6) by replacing $V^*(x)$ directly.

7. Review of Nonlinear MPC based on Nominal Models

The ultimate objective of a model predictive controller is to provide a *closed-loop feedback $u = \kappa_{mpc}(x)$* that regulates (1) to its target set (assumed here $x = 0$) in a fashion that is optimal with respect to the *infinite-time* problem (6), while enforcing pointwise constraints of the form $(x, u) \in \mathbb{X} \times \mathbb{U}$ in a constructive manner. However, rather than defining the map $\kappa_{mpc} : \mathbb{X} \to \mathbb{U}$ by solving a PDE of the form (8) (i.e thereby pre-computing knowledge of $\kappa_{mpc}(x)$ for *every* $x \in \mathbb{X}$), the model predictive control philosophy is to solve for, at time t, the control move $u = \kappa_{mpc}(x(t))$ for the *particular* value $x(t) \in \mathbb{X}$. This makes the online calculations inherently trajectory-based, and therefore closely tied to the results in Section 4 (with the caveat that the initial conditions are continuously referenced relative to current (t, x)). Since it is not practical to pose (online) trajectory-based calculations over an infinite prediction horizon $\tau \in [t, \infty)$, a truncated prediction $\tau \in [t, t+T]$ is used instead. The truncated tail of the integral in (6) is replaced by a (designer-specified) terminal penalty $W : \mathbb{X}_f \to \mathbb{R}_{\geq 0}$, defined over any local neighbourhood $\mathbb{X}_f \subset \mathbb{X}$ of the target $x = 0$. This results in a feedback of the form:

$$u = \kappa_{mpc}(x(t)) \triangleq u^*_{[t, t+T]}(t) \tag{11a}$$

where $u^*_{[t, t+T]}$ denotes the solution to the $x(t)$-dependent problem:

$$u^*_{[t, t+T]} \triangleq \arg \min_{u^p_{[t, t+T]}} \left(V_T(x(t), u^p_{[t, t+T]}) \triangleq \int_t^{t+T} L(x^p, u^p) \, d\tau + W(x^p(t+T)) \right) \tag{11b}$$

$$s.t. \quad \forall \tau \in [t, t+T] : \quad \frac{d}{d\tau} x^p = f(x^p, u^p), \quad x^p(t) = x(t) \tag{11c}$$

$$(x^p(\tau), u^p(\tau)) \in \mathbb{X} \times \mathbb{U} \tag{11d}$$

$$x^p(t+T) \in \mathbb{X}_f \tag{11e}$$

Clearly, if one could define $W(x) \equiv V^*(x)$ globally, then the feedback in (11) must satisfy $\kappa_{mpc}(\cdot) \equiv k_{DP}(\cdot)$. While $W(x) \equiv V^*(x)$ is generally unachievable, this motivates the selection of $W(x)$ as a CLF such that $W(x)$ is an inverse-optimal approximation of $V^*(x)$. A more precise characterization of the selection of $W(x)$ is the focus of the next section.

8. General Sufficient Conditions for Stability

A very general proof of the closed-loop stability of (11), which unifies a variety of earlier, more restrictive, results is presented[6] in the survey Mayne et al. (2000). This proof is based upon the following set of sufficient conditions for closed-loop stability:

Criterion 8.1. *The function $W : \mathbb{X}_f \to \mathbb{R}_{\geq 0}$ and set $\mathbb{X}_f$ are such that a local feedback $k_f : \mathbb{X}_f \to \mathbb{U}$ exists to satisfy the following conditions:*

C1) $0 \in \mathbb{X}_f \subseteq \mathbb{X}$, $\mathbb{X}_f$ closed   (i.e., state constraints satisfied in $\mathbb{X}_f$)

C2) $k_f(x) \in \mathbb{U}, \forall x \in \mathbb{X}_f$ (i.e., control constraints satisfied in $\mathbb{X}_f$)

C3) $\mathbb{X}_f$ is positively invariant for $\dot{x} = f(x, k_f(x))$.

C4) $L(x, k_f(x)) + \frac{\partial W}{\partial x} f(x, k_f(x)) \leq 0, \quad \forall x \in \mathbb{X}_f.$

Only existence, not knowledge, of $k_f(x)$ is assumed. Thus by comparison with (9), it can be seen that C4 essentially requires that $W(x)$ be a CLF over the (local) domain $\mathbb{X}_f$, in a manner consistent with the constraints.

In hindsight, it is nearly obvious that closed-loop stability can be reduced entirely to conditions placed upon only the terminal choices $W(\cdot)$ and $\mathbb{X}_f$. Viewing $V_T(x(t), u^*_{[t,t+T]})$ as a Lyapunov function candidate, it is clear from (3) that V_T contains "energy" in both the $\int L \, d\tau$ and terminal W terms. Energy dissipates from the front of the integral at a rate $L(x, u)$ as time t flows, and by the principle of optimality one could implement (11) on a *shrinking* horizon (i.e., $t + T$ constant), which would imply $\dot{V} = -L(x, u)$. In addition to this, C4 guarantees that the energy transfer from W to the integral (as the point $t + T$ recedes) will be non-increasing, and could even dissipate additional energy as well.

9. Robustness Considerations

As can be seen in Proposition 4.1, the presence of inequality constraints on the state variables poses a challenge for numerical solution of the optimal control problem in (11). While locating the times $\{t_i\}$ at which the active set changes can itself be a burdensome task, a significantly more challenging task is trying to guarantee that the tangency condition $N(x(t_{i+1})) = 0$ is met, which involves determining if x lies on (or crosses over) the critical surface beyond which this condition fails.

As highlighted in Grimm et al. (2004), this critical surface poses more than just a computational concern. Since both the cost function and the feedback $\kappa_{mpc}(x)$ are potentially discontinuous on this surface, there exists the potential for *arbitrarily small* disturbances (or other plant-model mismatch) to compromise closed-loop stability. This situation arises when the optimal solution $u^*_{[t,t+T]}$ in (11) switches between disconnected minimizers, potentially resulting in invariant limit cycles (for example, as a very low-cost minimizer alternates between being judged feasible/infeasible.)

A modification suggested in Grimm et al. (2004) to restore nominal robustness, similar to the idea in Marruedo et al. (2002), is to replace the constraint $x(\tau) \in \mathbb{X}$ of (11d) with one of the form $x(\tau) \in \mathbb{X}^o(\tau - t)$, where the function $\mathbb{X}^o : [0, T] \to \mathbb{X}$ satisfies $\mathbb{X}^o(0) = \mathbb{X}$, and the strict containment $\mathbb{X}^o(t_2) \subset \mathbb{X}^o(t_1)$, $t_1 < t_2$. The gradual relaxation of the constraint limit as future predictions move closer to current time provides a safety margin that helps to avoid constraint violation due to small disturbances.

---

[6] in the context of both continuous- and discrete-time frameworks

The issue of robustness to measurement error is addressed in Tuna et al. (2005). On one hand, nominal robustness to measurement noise of an MPC feedback was already established in Grimm et al. (2003) for discrete-time systems, and in Findeisen et al. (2003) for sampled-data implementations. However, Tuna et al. (2005) demonstrates that as the sampling frequency becomes arbitrarily fast, the margin of this robustness may approach zero. This stems from the fact that the feedback $\kappa_{mpc}(x)$ of (11) is inherently discontinuous in x if the indicated minimization is performed globally on a nonconvex surface, which by Coron & Rosier (1994); Hermes (1967) enables a fast measurement dither to generate flow in any direction contained in the convex hull of the discontinuous closed-loop vectorfield. In other words, additional attractors or unstable/infeasible modes can be introduced into the closed-loop behaviour by arbitrarily small measurement noise.

Although Tuna et al. (2005) deals specifically with situations of obstacle avoidance or stabilization to a target set containing disconnected points, other examples of problematic nonconvexities are depicted in Figure 1. In each of the scenarios depicted in Figure 1, measurement dithering could conceivably induce flow along the dashed trajectories, thereby resulting in either constraint violation or convergence to an undesired equilibrium.

Two different techniques were suggested in Tuna et al. (2005) for restoring robustness to the measurement error, both of which involve adding a hysteresis-type behaviour in the optimization to prevent arbitrary switching of the solution between separate minimizers (i.e., making the optimization behaviour more decisive).


Fig. 1. Examples of nonconvexities susceptible to measurement error

10. Robust MPC

10.1 Review of Nonlinear MPC for Uncertain Systems

While a vast majority of the robust-MPC literature has been developed within the framework of discrete-time systems[7], for consistency with the rest of this thesis most of the discussion will be based in terms of their continuous-time analogues. The uncertain system model is

---

[7] Presumably for numerical tractability, as well as providing a more intuitive link to game theory.

therefore described by the general form

$$\dot{x} = f(x, u, d) \tag{12}$$

where $d(t)$ represents any arbitrary $\mathcal{L}_\infty$-bounded disturbance signal, which takes point-wise[8] values $d \in \mathcal{D}$. Equivalently, (12) can be represented as the differential inclusion model $\dot{x} \in F(x, u) \triangleq f(x, u, \mathcal{D})$.

In the next two sections, we will discuss approaches for accounting explicitly for the disturbance in the online MPC calculations. We note that significant effort has also been directed towards various means of increasing the inherent robustness of the controller without requiring explicit online calculations. This includes the suggestion in Magni & Sepulchre (1997) (with a similar discrete-time idea in De Nicolao et al. (1996)) to use a modified stage cost $\overline{L}(x, u) \triangleq L(x, u) + \langle \nabla_x V_T^*(x), f(x, u) \rangle$ to increase the robustness of a nominal-model implementation, or the suggestion in Kouvaritakis et al. (2000) to use an prestabilizer, optimized offline, of the form $u = Kx + v$ to reduced online computational burden. Ultimately, these approaches can be considered encompassed by the banner of nominal-model implementation.

10.1.1 Explicit robust MPC using Open-loop Models

As seen in the previous chapters, essentially all MPC approaches depend critically upon the Principle of Optimality (Def 3.1) to establish a proof of stability. This argument depends inherently upon the assumption that the predicted trajectory $x_{[t, t+T]}^p$ is an invariant set under open-loop implementation of the corresponding $u_{[t, t+T]}^p$; i.e., that the prediction model is "perfect". Since this is no longer the case in the presence of plant-model mismatch, it becomes necessary to associate with $u_{[t, t+T]}^p$ a *cone* of trajectories $\{x_{[t, t+T]}^p\}_\mathcal{D}$ emanating from $x(t)$, as generated by (12).

Not surprisingly, establishing stability requires a strengthening of the conditions imposed on the selection of the terminal cost W and domain $\mathbb{X}_f$. As such, W and $\mathbb{X}_f$ are assumed to satisfy Criterion (8.1), but with the revised conditions:

C3a) $\mathbb{X}_f$ is strongly positively invariant for $\dot{x} \in f(x, k_f(x), \mathcal{D})$.

C4a) $L(x, k_f(x)) + \frac{\partial W}{\partial x} f(x, k_f(x), d) \leq 0, \quad \forall (x, d) \in \mathbb{X}_f \times \mathcal{D}$.

While the original C4 had the interpretation of requiring W to be a CLF for the nominal system, so the revised C4a can be interpreted to imply that W should be a robust-CLF like those developed in Freeman & Kokotović (1996b).

Given such an appropriately defined pair $(W, \mathbb{X}_f)$, the model predictive controller explicitly considers all trajectories $\{x_{[t, t+T]}^p\}_\mathcal{D}$ by posing the modified problem

$$u = \kappa_{mpc}(x(t)) \triangleq u_{[t, t+T]}^*(t) \tag{13a}$$

where the trajectory $u_{[t, t+T]}^*$ denotes the solution to

$$u_{[t, t+T]}^* \triangleq \arg \min_{\substack{u_{[t, t+T]}^p \\ T \in [0, T_{max}]}} \left(\max_{d_{[t, t+T]} \in \mathcal{D}} V_T(x(t), u_{[t, t+T]}^p, d_{[t, t+T]}) \right) \tag{13b}$$

---

[8] The abuse of notation $d_{[t_1, t_2]} \in \mathcal{D}$ is likewise interpreted pointwise

The function $V_T(x(t), u^p_{[t, t+T]}, d_{[t, t+T]})$ appearing in (13) is as defined in (11), but with (11c) replaced by (12). Variations of this type of design are given in Chen et al. (1997); Lee & Yu (1997); Mayne (1995); Michalska & Mayne (1993); Ramirez et al. (2002), differing predominantly in the manner by which they select $W(\cdot)$ and $\mathbb{X}_f$.

If one interprets the word "optimal" in Definition 3.1 in terms of the worst-case trajectory in the optimal cone $\{x^p_{[t, t+T]}\}^*_{\mathcal{D}}$, then at time $\tau \in [t, t+T]$ there are only two possibilities:

- the actual $x_{[t, \tau]}$ matches the subarc from a worst-case element of $\{x^p_{[t, t+T]}\}^*_{\mathcal{D}}$, in which case the Principle of Optimality holds as stated.

- the actual $x_{[t, \tau]}$ matches the subarc from an element in $\{x^p_{[t, t+T]}\}^*_{\mathcal{D}}$ which was *not* the worst case, so implementing the remaining $u^*_{[\tau, t+T]}$ will achieve overall less cost than the worst-case estimate at time t.

One will note however, that the bound guaranteed by the principle of optimality applies only to the remaining subarc $[\tau, t+T]$, and says nothing about the ability to extend the horizon. For the nominal-model results of Chapter 7, the ability to extend the horizon followed from C4) of Criterion (8.1). In the present case, C4a) guarantees that for *each* terminal value $\{x^p_{[t, t+T]}(t+T)\}^*_{\mathcal{D}}$ there exists a value of u rendering W decreasing, but not necessarily a single such value satisfying C4a) for *every* $\{x^p_{[t, t+T]}(t+T)\}^*_{\mathcal{D}}$. Hence, receding of the horizon can only occur at the discretion of the optimizer. In the *worst* case, T could contract (i.e., $t+T$ remains fixed) until eventually $T = 0$, at which point $\{x^p_{[t, t+T]}(t+T)\}^*_{\mathcal{D}} \equiv x(t)$, and therefore by C4a) an appropriate extension of the "trajectory" $u^*_{[t,t]}$ exists.

Although it is not an explicit min-max type result, the approach in Marruedo et al. (2002) makes use of global Lipschitz constants to determine a bound on the the worst-case distance between a solution of the uncertain model (12), and that of the underlying nominal model estimate. This Lipschitz-based uncertainty cone expands at the fastest-possible rate, necessarily containing the actual uncertainty cone $\{x^p_{[t, t+T]}\}_{\mathcal{D}}$. Although ultimately just a nominal-model approach, it is relevant to note that it can be viewed as replacing the "max" in (13) with a simple worst-case upper bound.

Finally, we note that many similar results Cannon & Kouvaritakis (2005); Kothare et al. (1996) in the linear robust-MPC literature are relevant, since nonlinear dynamics can often be approximated using uncertain linear models. In particular, linear systems with polytopic descriptions of uncertainty are one of the few classes that can be realistically solved numerically, since the calculations reduce to simply evaluating each node of the polytope.

10.1.2 Explicit robust MPC using Feedback Models

Given that robust control design is closely tied to game theory, one can envision (13) as representing a player's decision-making process throughout the evolution of a strategic game. However, it is unlikely that a player even moderately-skilled at such a game would restrict themselves to preparing only a single *sequence of moves* to be executed in the future. Instead, a skilled player is more likely to prepare a *strategy* for future game-play, consisting of several "backup plans" contingent upon future responses of their adversary.

To be as least-conservative as possible, an ideal (in a worst-case sense) decision-making process would more properly resemble

$$u = \kappa_{mpc}(x(t)) \triangleq u^*_t \tag{14a}$$

where $u_t^* \in \mathbb{R}^m$ is the constant value satisfying

$$u_t^* \triangleq \arg\min_{u_t} \left(\max_{d_{[t,t+T]} \in \mathcal{D}} \min_{u_{[t,t+T]}^p \in \mathcal{U}(u_t)} V_T(x(t), u_{[t,t+T]}^p, d_{[t,t+T]}) \right) \tag{14b}$$

with the definition $\mathcal{U}(u_t) \triangleq \{u_{[t,t+T]}^p \mid u^p(t) = u_t\}$. Clearly, the "least conservative" property follows from the fact that a separate response is optimized for every possible sequence the adversary could play. This is analogous to the philosophy in Scokaert & Mayne (1998), for system $x^+ = Ax + Bu + d$, in which polytopic $\mathcal{D}$ allows the max to be reduced to selecting the worst index from a finitely-indexed collection of responses; this equivalently replaces the innermost minimization with an augmented search in the outermost loop over *all* input responses in the collection.

While (14) is useful as a definition, a more useful (equivalent) representation involves minimizing over *feedback policies* $k : [t, t+T] \times \mathbb{X} \to \mathbb{U}$ rather than trajectories:

$$u = \kappa_{mpc}(x(t)) \triangleq k^*(t, x(t)) \tag{15a}$$

$$k^*(\cdot, \cdot) \triangleq \arg\min_{k(\cdot, \cdot)} \max_{d_{[t,t+T]} \in \mathcal{D}} \left(V_T(x(t), k(\cdot, \cdot), d_{[t,t+T]}) \right) \tag{15b}$$

$$V_T(x(t), k(\cdot, \cdot), d_{[t,t+T]}) \triangleq \int_t^{t+T} L(x^p, k(\tau, x^p(\tau))) \, d\tau + W(x^p(t+T)) \tag{15c}$$

$$\text{s.t. } \forall \tau \in [t, t+T]: \qquad \frac{d}{d\tau} x^p = f(x^p, k(\tau, x^p(\tau)), d), \quad x^p(t) = x(t) \tag{15d}$$

$$(x^p(\tau), k(\tau, x^p(\tau))) \in \mathbb{X} \times \mathbb{U} \tag{15e}$$

$$x^p(t+T) \in \mathbb{X}_f \tag{15f}$$

There is a recursive-like elegance to (15), in that $\kappa_{mpc}(x)$ is essentially defined as a search over future candidates of itself. Whereas (14) explicitly involves *optimization-based* future feedbacks, the search in (15) can actually be (suboptimally) restricted to *any* arbitrary sub-class of feedbacks $k : [t, t+T] \times \mathbb{X} \to \mathbb{U}$. For example, this type of approach first appeared in Kothare et al. (1996); Lee & Yu (1997); Mayne (1995), where the cost functional was minimized by restricting the search to the class of linear feedback $u = Kx$ (or $u = K(t)x$).

The error cone $\{x_{[t,t+T]}^p\}_{\mathcal{D}}^*$ associated with (15) is typically *much* less conservative than that of (13). This is due to the fact that (15d) accounts for future disturbance attenuation resulting from $k(\tau, x^p(\tau))$, an effect ignored in the open-loop predictions of (13). In the case of (14) and (15) it is no longer necessary to include T as an optimization variable, since by condition C4a one can now envision extending the horizon by appending an increment $k(T+\delta t, \cdot) = k_f(\cdot)$. This notion of feedback MPC has been applied in Magni et al. (2003; 2001) to solve $\mathcal{H}_\infty$ disturbance attenuation problems. This approach avoids the need to solve a difficult Hamilton-Jacobi-Isaacs equation, by combining a specially-selected stage cost $L(x, u)$ with a local HJI approximation $W(x)$ (designed generally by solving an $\mathcal{H}_\infty$ problem for the linearized system). An alternative perspective of the implementation of (15) is developed in Langson et al. (2004), with particular focus on obstacle-avoidance in Raković & Mayne (2005). In this work, a set-invariance philosophy is used to propagate the uncertainty cone $\{x_{[t,t+T]}^p\}_{\mathcal{D}}$ for (15d) in the form of a control-invariant tube. This enables the use of efficient methods for constructing control invariant sets based on approximations such as polytopes or ellipsoids.

11. Adaptive Approaches to MPC

The sectionr will be focused on the more typical role of adaptation as a means of coping with uncertainties in the system model. A standard implementation of model predictive control using a nominal model of the system dynamics can, with slight modification, exhibit nominal robustness to disturbances and modelling error. However in practical situations, the system model is only approximately known, so a guarantee of robustness which covers only "sufficiently small" errors may be unacceptable. In order to achieve a more solid robustness guarantee, it becomes necessary to account (either explicitly, or implicitly) for *all* possible trajectories which could be realized by the uncertain system, in order to guarantee feasible stability. The obvious numerical complexity of this task has resulted in an array of different control approaches, which lie at various locations on the spectrum between simple, conservative approximations versus complex, high-performance calculations. Ultimately, selecting an appropriate approach involves assessing, for the particular system in question, what is an acceptable balance between computational requirements and closed-loop performance.

Despite the fact that the ability to adjust to changing process conditions was one of the earliest industrial motivators for developing predictive control techniques, the progress in this area has been negligible. The small amount of progress that has been made is restricted to systems which do not involve constraints on the state, and which are affine in the unknown parameters. We will briefly describe two such results.

11.1 Certainty-equivalence Implementation

The result in Mayne & Michalska (1993) implements a certainty equivalence nominal-model[9] MPC feedback of the form $u(t) = \kappa_{mpc}(x(t), \hat{\theta}(t))$, to stabilize the uncertain system

$$\dot{x} = f(x, u, \theta) \triangleq f_0(x, u) + g(x, u)\theta \tag{16}$$

subject to an input constraint $u \in \mathbb{U}$. The vector $\theta \in \mathbb{R}^p$ represents a set of unknown constant parameters, with $\hat{\theta} \in \mathbb{R}^p$ denoting an identifier. Certainty equivalence implies that the nominal prediction model (11c) is of the same form as (16), but with $\hat{\theta}$ used in place of θ.

At any time $t \geq 0$, the identifier $\hat{\theta}(t)$ is defined to be a (min-norm) solution of

$$\int_0^t g(x(s), u(s))^T \Big(\dot{x}(s) - f_0(x(s), u(s))\Big) ds = \int_0^t g(x(s), u(s))^T g(x(s), u(s)) ds \, \hat{\theta} \tag{17}$$

which is solved over the window of all past history, under the assumption that $\dot{x}$ is measured (or computable). If necessary, an additional search is performed along the nullspace of $\int_0^t g(x, u)^T g(x, u) ds$ in order to guarantee $\hat{\theta}(t)$ yields a controllable certainty-equivalence model (since (17) is controllable by assumption).

The final result simply shows that there must exist a time $0 < t_a < \infty$ such that the regressor $\int_0^t g(x, u)^T g(x, u) ds$ achieves full rank, and thus $\hat{\theta}(t) \equiv \theta$ for all $t \geq t_a$. However, it is only by assumption that the state $x(t)$ does not escape the stabilizable region during the identification phase $t \in [0, t_a]$; moreover, there is no mechanism to decrease t_a in any way, such as by injecting excitation.

---

[9] Since this result arose early in the development of nonlinear MPC, it happens to be based upon a terminal-constrained controller (i.e., $\mathbb{X}_f \equiv \{0\}$); however, this is not critical to the adaptation.

11.1.1 Stability-Enforced Approach

One of the early stability results for nominal-model MPC in (Primbs (1999); Primbs et al. (2000)) involved the use of a global CLF $V(x)$ instead of a terminal penalty. Stability was enforced by constraining the optimization such that $V(x)$ is decreasing, and performance achieved by requiring the predicted cost to be less than that accumulated by simulation of pointwise min-norm control.

This idea was extended in Adetola & Guay (2004) to stabilize unconstrained systems of the form

$$\dot{x} = f(x, u, \theta) \triangleq f_0(x) + g_\theta(x)\theta + g_u(x)u \tag{18}$$

Using ideas from robust stabilization, it is assumed that a global ISS-CLF[10] is known for the nominal system. Constraining $V(x)$ to decrease ensures convergence to a neighbourhood of the origin, which gradually contracts as the identification proceeds. Of course, the restrictiveness of this approach lies in the assumption that $V(x)$ is known.

12. An Adaptive Approach to Robust MPC

Both the theoretical and practical merits of model-based predictive control strategies for nonlinear systems are well established, as reviewed in Chapter 7. To date, the vast majority of implementations involve an "accurate model" assumption, in which the control action is computed on the basis of predictions generated by an approximate nominal process model, and implemented (un-altered) on the actual process. In other words, the effects of plant-model mismatch are completely ignored in the control calculation, and closed-loop stability hinges upon the critical assumption that the nominal model is a "sufficiently close" approximation of the actual plant. Clearly, this approach is only acceptable for processes whose dynamics can be modelled a-priori to within a high degree of precision.

For systems whose true dynamics can only be approximated to within a large margin of uncertainty, it becomes necessary to directly account for the plant-model mismatch. To date, the most general and rigourous means for doing this involves explicitly accounting for the error in the online calculation, using the robust-MPC approaches discussed in Section 10.1. While the theoretical foundations and guarantees of stability for these tools are well established, it remains problematic in most cases to find an appropriate approach yielding a satisfactory balance between computational complexity, and conservatism of the error calculations. For example, the framework of min-max feedback-MPC Magni et al. (2003); Scokaert & Mayne (1998) provides the least-conservative control by accounting for the effects of future feedback actions, but is in most cases computationally intractable. In contrast, computationally simple approaches such as the openloop method of Marruedo et al. (2002) yield such conservatively-large error estimates, that a feasible solution to the optimal control problem often fails to exist. For systems involving primarily *static* uncertainties, expressible in the form of unknown (constant) model parameters $\theta \in \Theta \subset \mathbb{R}^p$, it would be more desirable to approach the problem in the framework of adaptive control than that of robust control. Ideally, an adaptive mechanism enables the controller to improve its performance over time by employing a process model which asymptotically approaches that of the true system. Within the context of predictive control, however, the transient effects of parametric estimation error have proven problematic

---

[10] i.e., a CLF guaranteeing robust stabilization to a neighbourhood of the origin, where the size of the neighbourhood scales with the $\mathcal{L}_\infty$ bound of the disturbance signal

towards developing anything beyond the limited results discussed in Section 11. In short, the development of a general "robust adaptive-MPC" remains at present an open problem.

In the following, we make no attempt to construct such a "robust adaptive" controller; instead we propose an approach more properly referred to as "adaptive robust" control. The approach differs from typical adaptive control techniques, in that the adaptation mechanism does not directly involve a parameter identifier $\hat{\theta} \in \mathbb{R}^p$. Instead, a set-valued description of the parametric uncertainty, Θ, is adapted online by an identification mechanism. By gradually eliminating values from Θ that are identified as being inconsistent with the observed trajectories, Θ gradually contracts upon θ in a nested fashion. By virtue of this nested evolution of Θ, it is clear that an adaptive feedback structure of the form in Figure 2 would retain the stability properties of any underlying robust control design.
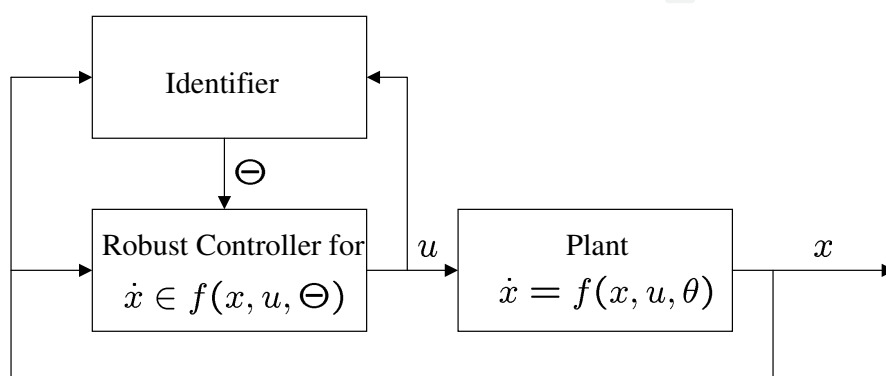

Fig. 2. Adaptive robust feedback structure

The idea of arranging an identifier and robust controller in the configuration of Figure 2 is itself not entirely new. For example the robust control design of Corless & Leitmann (1981), appropriate for nonlinear systems affine in u whose disturbances are bounded and satisfy the so-called "matching condition", has been used by various authors Brogliato & Neto (1995); Corless & Leitmann (1981); Tang (1996) in conjunction with different identifier designs for estimating the disturbance bound resulting from parametric uncertainty. A similar concept for linear systems is given in Kim & Han (2004).

However, to the best of our knowledge this idea has not been well explored in the situation where the underlying robust controller is designed by robust-MPC methods. The advantage of such an approach is that one could then potentially imbed an internal model of the identification mechanism into the predictive controller, as shown in Figure 3. In doing so the effects of future identification are accounted for within the optimal control problem, the benefits of which are discussed in the next section.

13. A Minimally-Conservative Perspective

13.1 Problem Description

The problem of interest is to achieve robust regulation, by means of state-feedback, of the system state to some compact target set $\Sigma_x^o \in \mathbb{R}^n$. Optimality of the resulting trajectories are measured with respect to the accumulation of some instantaneous penalty (i.e., stage cost) $L(x, u) \geq 0$, which may or may not have physical significance. Furthermore, the state and input trajectories are required to obey pointwise constraints $(x, u) \in \mathbb{X} \times \mathbb{U} \subseteq \mathbb{R}^n \times \mathbb{R}^m$.
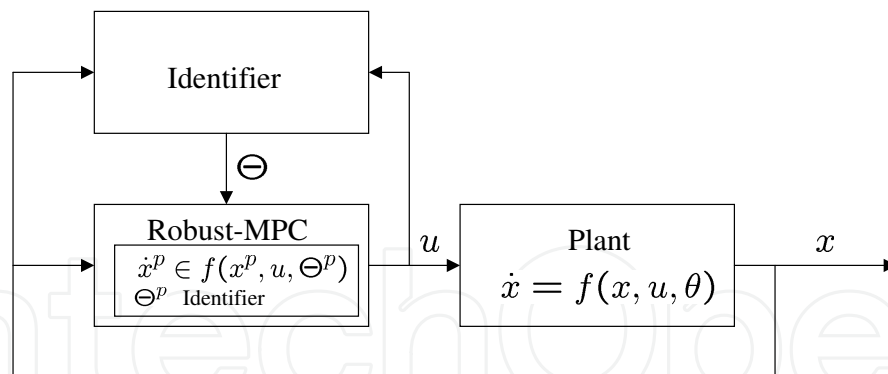
Fig. 3. Adaptive robust MPC structure

It is assumed that the system dynamics are not fully known, with uncertainty stemming from both unmodelled static nonlinearities as well as additional exogenous inputs. As such, the dynamics are assumed to be of the general form

$$\dot{x} = f(x, u, \theta, d(t)) \tag{19}$$

where f is a locally Lipschitz vector function of state $x \in \mathbb{R}^n$, control input $u \in \mathbb{R}^m$, disturbance input $d \in \mathbb{R}^d$, and constant parameters $\theta \in \mathbb{R}^p$. The entries of θ may represent physically meaningful model parameters (whose values are not exactly known *a-priori*), or alternatively they could be parameters associated with any (finite) set of universal basis functions used to approximate unknown nonlinearities. The disturbance $d(t)$ represents the combined effects of actual exogenous inputs, neglected system states, or static nonlinearities lying outside the span of θ (such as the truncation error resulting from using a finite basis).

Assumption 13.1. $\theta \in \Theta^o$, where Θ^o is a known compact subset of $\mathbb{R}^p$.

Assumption 13.2. $d(\cdot) \in \mathcal{D}_\infty$, where $\mathcal{D}_\infty$ is the set of all right-continuous $\mathcal{L}^\infty$-bounded functions $d : \mathbb{R} \to \mathcal{D}$; i.e., composed of continuous subarcs $d_{[a,b)}$, and satisfying $d(\tau) \in \mathcal{D}$, $\forall \tau \in \mathbb{R}$, with $\mathcal{D} \subset \mathbb{R}^d$ a compact vectorspace.   ∎

Unlike much of the robust or adaptive MPC literature, we do not necessarily assume exact knowledge of the system equilibrium manifold, or its stabilizing equilibrium control map. Instead, we make the following (weaker) set of assumptions:

Assumption 13.3. Letting $\Sigma_u^o \subseteq \mathbb{U}$ be a chosen compact set, assume that $L : \mathbb{X} \times \mathbb{U} \to \mathbb{R}_{\geq 0}$ is continuous, $L(\Sigma_x^o, \Sigma_u^o) \equiv 0$, and $L(x, u) \geq \underline{\gamma}_L \left(\|(x, u)\|_{\Sigma_x^o \times \Sigma_u^o} \right)$, $\underline{\gamma}_L \in \mathcal{K}_\infty$. As well, assume that

$$\min_{(u, \theta, d) \in \mathbb{U} \times \Theta^o \times \mathcal{D}} \left(\frac{L(x, u)}{\|f(x, u, \theta, d)\|} \right) \geq \frac{c_2}{\|x\|_{\Sigma_x^o}} \qquad \forall x \in \mathbb{X} \setminus B(\Sigma_x^o, c_1) \tag{20}$$

∎

Definition 13.4. For each $\Theta \subseteq \Theta^o$, let $\Sigma_x(\Theta) \subseteq \Sigma_x^o$ denote the maximal (strongly) control-invariant subset for the differential inclusion $\dot{x} \in f(x, u, \Theta, \mathcal{D})$, using only controls $u \in \Sigma_u^o$.

Assumption 13.5. There exists a constant $N_\Sigma < \infty$, and a finite cover of Θ^o (not necessarily unique), denoted $\{\Theta\}^\Sigma$, such that

    i. the collection $\{\mathring{\Theta}\}^{\Sigma}$ is an open cover for the interior $\mathring{\Theta}^{o}$.

    ii. $\Theta \in \{\Theta\}^{\Sigma}$ implies $\Sigma_{x}(\Theta) \neq \varnothing$.

    iii. $\{\Theta\}^{\Sigma}$ contains at most N_{Σ} elements.   ■

The most important requirement of Assumption 13.3 is that, since the exact location (in $\mathbb{R}^{n} \times \mathbb{R}^{m}$) of the equilibrium[11] manifold is not known *a-priori*, $L(x, u)$ must be identically zero on the entire region of equilibrium candidates $\Sigma_{x}^{o} \times \Sigma_{u}^{o}$. One example of how to construct such a function would be to define $L(x, u) = \rho(x, u)\overline{L}(x, u)$, where $\overline{L}(x, u)$ is an arbitrary penalty satisfying $(x, u) \notin \Sigma_{x}^{o} \times \Sigma_{u}^{o} \implies \overline{L}(x, u) > 0$, and $\rho(x, u)$ is a smoothed indicator function of the form

$$\rho(x, u) = \begin{cases} 0 & (x, u) \in \Sigma_{x}^{o} \times \Sigma_{u}^{o} \\ \frac{\|(x,u)\|_{\Sigma_{x}^{o} \times \Sigma_{u}^{o}}}{\delta_{\rho}} & 0 < \|(x, u)\|_{\Sigma_{x}^{o} \times \Sigma_{u}^{o}} < \delta_{\rho} \\ 1 & \|(x, u)\|_{\Sigma_{x}^{o} \times \Sigma_{u}^{o}} \geq \delta_{\rho} \end{cases} \tag{21}$$

The restriction that $L(x, u)$ is strictly positive definite with respect to $\Sigma_{x}^{o} \times \Sigma_{u}^{o}$ is made for convenience, and could be relaxed to positive semi-definite using an approach similar to that of Grimm et al. (2005) as long as $L(x, u)$ satisfies an appropriate detectability assumption (i.e., as long as it is guaranteed that all trajectories remaining in $\{x \mid \exists u \ s.t. \ L(x, u) = 0\}$ must asymptotically approach $\Sigma_{x}^{o} \times \Sigma_{u}^{o}$).

The first implication of Assumption 13.5 is that for any $\theta \in \Theta^{o}$, the target Σ_{x}^{o} contains a stabilizable "equilibrium" $\Sigma(\theta)$ such that the regulation problem is well-posed. Secondly, the openness of the covering in Assumption 13.5 implies a type of "local-ISS" property of these equilibria with respect to perturbations in small neighbourhoods Θ of θ. This property ensures that the target is stabilizable given "sufficiently close" identification of the unknown θ, such that the adaptive controller design is tractable.

13.2 Adaptive Robust Controller Design Framework
13.2.1 Adaptation of Parametric Uncertainty Sets

Unlike standard approaches to adaptive control, this work does not involve explicitly generating a parameter estimator $\hat{\theta}$ for the unknown θ. Instead, the parametric uncertainty set Θ^{o} is adapted to gradually eliminate sets which do not contain θ. To this end, we define the infimal uncertainty set

$$\mathbb{Z}(\Theta, x_{[a,b]}, u_{[a,b]}) \triangleq \{\theta \in \Theta \mid \dot{x}(\tau) \in f(x(\tau), u(\tau), \theta, \mathcal{D}), \ \forall \tau \in [a, b]\} \tag{22}$$

By definition, $\mathbb{Z}$ represents the best-case performance that could be achieved by any identifier, given a set of data generated by (19), and a prior uncertainty bound Θ. Since exact online calculation of (22) is generally impractical, we assume that the set $\mathbb{Z}$ is approximated online using an arbitrary estimator Ψ. This estimator must be chosen to satisfy the following conditions.

Criterion 13.6. $\Psi(\cdot, \cdot, \cdot)$ *is designed such that for $a \leq b \leq c$, and for any $\Theta \subseteq \Theta^{o}$,*

    C13.6.1 $\mathbb{Z} \subseteq \Psi$

    C13.6.2 $\Psi(\Theta, \cdot, \cdot) \subseteq \Theta$, and closed.

---

[11] we use the word "equilibrium" loosely in the sense of control-invariant subsets of the target Σ_{x}^{o}, which need not be actual equilibrium *points* in the traditional sense

C13.6.3 $\Psi(\Theta_1, x_{[a,b]}, u_{[a,b]}) \subseteq \Psi(\Theta_2, x_{[a,b]}, u_{[a,b]}), for \Theta_1 \subseteq \Theta_2 \subseteq \Theta^o$

C13.6.4 $\Psi(\Theta, x_{[a,b]}, u_{[a,b]}) \supseteq \Psi(\Theta, x_{[a,c]}, u_{[a,c]})$

C13.6.5 $\Psi(\Theta, x_{[a,c]}, u_{[a,c]}) \equiv \Psi(\Psi(\Theta, x_{[a,b]}, u_{[a,b]}), x_{[b,c]}, u_{[b,c]})$

The set Ψ represents an approximation of $\mathbb{Z}$ in two ways. First, both Θ^o and Ψ can be restricted *a-priori* to any class of finitely-parameterized sets, such as linear polytopes, quadratic balls, etc. Second, contrary to the actual definition of (22), Ψ can be computed by *removing* values from Θ^o as they are determined to violate the differential inclusion model. As such, the search for infeasible values can be terminated at any time without violating C13.6.

The closed loop dynamics of (19) then take the form

$$\dot{x} = f(x, \kappa_{mpc}(x, \Theta(t)), \theta, d(t)), \quad x(t_0) = x_0 \tag{23a}$$

$$\Theta(t) = \Psi(\Theta^o, x_{[t_0,\ t]}, u_{[t_0,\ t]}) \tag{23b}$$

where $\kappa_{mpc}(x, \Theta)$ represents the MPC feedback policy, detailed in Section 13.2.2. In practice, the (set-valued) controller state Θ could be generated using an update law $\dot{\Theta}$ designed to gradually contract the set (satisfying C13.6). However, the given statement of (23b) is more general, as it allows for $\Theta(t)$ to evolve discontinuously in time, as may happen for example when the sign of a parameter can suddenly be conclusively determined.

13.2.2 Feedback-MPC framework

In the context of min-max robust MPC, it is well known that feedback-MPC, because of its ability to account for the effects of future feedback decisions on disturbance attenuation, provides significantly less conservative performance than standard open-loop MPC implementations. In the following, the same principle is extended to incorporate the effects of future parameter adaptation.

In typical feedback-MPC fashion, the receding horizon control law in (23) is defined by minimizing over *feedback policies* $\kappa : \mathbb{R}_{\geq 0} \times \mathbb{R}^n \times \overline{\text{cov}}\{\Theta^o\} \to \mathbb{R}^m$ as

$$u = \kappa_{mpc}(x, \Theta) \triangleq \kappa^*(0, x, \Theta) \tag{24a}$$

$$\kappa^* \triangleq \arg \min_{\kappa(\cdot,\cdot,\cdot)} J(x, \Theta, \kappa) \tag{24b}$$

where $J(x, \Theta, \kappa)$ is the (worst-case) cost associated with the optimal control problem:

$$J(x, \Theta, \kappa) \triangleq \max_{\substack{\theta \in \Theta \\ d(\cdot) \in \mathcal{D}_\infty}} \int_0^T L(x^p, u^p) d\tau + W(x_f^p, \hat{\Theta}_f) \tag{25a}$$

$$s.t. \quad \forall \tau \in [0, T]$$

$$\frac{d}{d\tau} x^p = f(x^p, u^p, \theta, d), \ x^p(0) = x \tag{25b}$$

$$\hat{\Theta}(\tau) = \Psi_p(\Theta(t), x^p_{[0,\tau]}, u^p_{[0,\tau]}) \tag{25c}$$

$$x^p(\tau) \in \mathbb{X} \tag{25d}$$

$$u^p(\tau) \triangleq \kappa(\tau, x^p(\tau), \hat{\Theta}(\tau)) \in \mathbb{U} \tag{25e}$$

$$x_f^p \triangleq x^p(T) \in \mathbb{X}_f(\hat{\Theta}_f) \tag{25f}$$

$$\hat{\Theta}_f \triangleq \Psi_f(\Theta(t), x^p_{[0,T]}, u^p_{[0,T]}) \tag{25g}$$

Throughout the remainder, we denote the optimal cost $J^*(x, \Theta) \triangleq J(x, \Theta, \kappa^*)$, and furthermore we drop the explicit constraints (25d)-(25f) by assuming the definitions of L and W have been extended as follows:

$$L(x, u) = \begin{cases} L(x, u) < \infty & (x, u) \in \mathbb{X} \times \mathbb{U} \\ +\infty & \text{otherwise} \end{cases} \tag{26a}$$

$$W(x, \Theta) = \begin{cases} W(x, \Theta) < \infty & x \in \mathbb{X}_f(\Theta) \\ +\infty & \text{otherwise} \end{cases} \tag{26b}$$

The parameter identifiers Ψ_p and Ψ_f in (25) represent internal model approximations of the actual identifier Ψ, and must satisfy both C13.6 as well as the following criterion:

Criterion 13.7. *For identical arguments, $\mathbb{Z} \subseteq \Psi \subseteq \Psi_f \subseteq \Psi_p$.*

Remark 13.8. *We distinguish between different identifiers to emphasize that, depending on the frequency at which calculations are called, differing levels of accuracy can be applied to the identification calculations. The ordering in Criterion 13.7 is required for stability, and implies that identifiers existing within faster timescales provide more conservative approximations of the uncertainty set.*

There are two important characteristics which distinguish (25) from a standard (non-adaptive) feedback-MPC approach. First, the future evolution of $\hat{\Theta}$ in (25c) is fed back into both (25b) and (25e). The benefits of this feedback are analogous to those of adding state-feedback into the MPC calculation; the resulting cone of possible trajectories $x^p(\cdot)$ is narrowed by accounting for the effects of future adaptation on disturbance attenuation, resulting in less conservative worst-case predictions.

The second distinction is that both W and $\mathbb{X}_f$ are parameterized as functions of $\hat{\Theta}_f$, which reduces the conservatism of the terminal cost. Since the terminal penalty W has the interpretation of the "worst-case cost-to-go", it stands to reason that W should decrease with decreased parametric uncertainty. In addition, the domain $\mathbb{X}_f$ would be expected to enlarge with decreased parametric uncertainty, which in some situations could mean that a stabilizing CLF-pair $(W(x, \Theta), \mathbb{X}_f(\Theta))$ can be constructed even when no such CLF exists for the original uncertainty Θ^o. This effect is discussed in greater depth in Section 14.1.1.

13.2.3 Generalized Terminal Conditions

To guide the selection of $W(x_f, \hat{\Theta}_f)$ and $\mathbb{X}_f(\hat{\Theta}_f)$ in (25), it is important to outline (sufficient) conditions under which (23)-(25) can guarantee stabilization to the target Σ_x^o. The statement given here is extended from the set of such conditions for robust MPC from Mayne et al. (2000) that was outlined in Sections 8 and 10.1.1.

For reasons that are explained later in Section 14.1.1, it is useful to present these conditions in a more general context in which $W(\cdot, \Theta)$ is allowed to be $\mathcal{LS}$-continuous with respect to x, as may occur if W is generated by a switching mechanism. This adds little additional complexity to the analysis, since (25) is already discontinuous due to constraints.

Criterion 13.9. *The set-valued terminal constraint function $\mathbb{X}_f : \overline{cov}\{\Theta^o\} \to \overline{cov}\{\mathbb{X}\}$ and terminal penalty function $W : \mathbb{R}^n \times \overline{cov}\{\Theta^o\} \to [0, +\infty]$ are such that for each $\Theta \in \overline{cov}\{\Theta^o\}$, there exists $k_f(\cdot, \Theta) : \mathbb{X}_f \to \mathbb{U}$ satisfying*

    C13.9.1 $\mathbb{X}_f(\Theta) \neq \varnothing$ implies that $\Sigma_x^o \cap \mathbb{X}_f(\Theta) \neq \varnothing$, and $\mathbb{X}_f(\Theta) \subseteq \mathbb{X}$ is closed

    C13.9.2 $W(\cdot, \Theta)$ is $\mathcal{LS}$-continuous with respect to $x \in \mathbb{R}^n$

C13.9*.3* $k_f(x, \Theta) \in \mathbb{U}$, *for all* $x \in \mathbb{X}_f(\Theta)$.

C13.9*.4* $\mathbb{X}_f(\Theta)$ *and* $\Sigma_x(\Theta) \subseteq \left\{ \Sigma_x^o \cap \mathbb{X}_f(\Theta) \right\}$ *are both strongly positively invariant with respect to the differential inclusion* $\dot{x} \in f(x, k_f(x, \Theta), \Theta, \mathcal{D})$.

C13.9*.5* $\forall x \in \mathbb{X}_f(\Theta)$, *and denoting* $\mathcal{F} \triangleq f(x, k_f(x, \Theta), \Theta, \mathcal{D})$,

$$\max_{f \in \mathcal{F}} \left(L(x, k_f(x, \Theta)) + \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \left(\frac{W(x + \delta v, \Theta) - W(x, \Theta)}{\delta} \right) \right) \leq 0$$

∎

Although condition C13.9.5 is expressed in a slightly non-standard form, it embodies the standard interpretation that W must be decreasing by at least an amount $-L(x, k_f(x, \Theta))$ along all vectorfields in the closed-loop differential inclusion $\mathcal{F}$; i.e., $W(x, \Theta)$ is a robust-CLF (in an appropriate non-smooth sense) on the domain $\mathbb{X}_f(\Theta)$. Lyapunov stability involving $\mathcal{LS}$-continuous functions is thoroughly studied in Clarke et al. (1998), and provides a meaningful sense in which W can be considered a "robust-CLF" despite its discontinuous nature.

It is important to note that for the purposes of Criterion 13.9, $W(x, \Theta)$ and $\mathbb{X}_f(\Theta)$ are parameterized by the set Θ, but the criterion imposes no restrictions on their functional dependence with respect to the Θ argument. This Θ-dependence is required to satisfy the following criteria:

Criterion 13.10. *For any* $\Theta_1, \Theta_2 \in \overline{cov}\left\{\Theta^o\right\}$ *such that* $\Theta_1 \subseteq \Theta_2$,

  C13.10*.1* $\mathbb{X}_f(\Theta_2) \subseteq \mathbb{X}_f(\Theta_1)$

  C13.10*.2* $W(x, \Theta_1) \leq W(x, \Theta_2)$, $\forall x \in \mathbb{X}_f(\Theta_2)$

∎

Designing W and $\mathbb{X}_f$ as functions of Θ satisfying Criteria 13.9 and 13.10 may appear prohibitively complex; however, the task is greatly simplified by noting that neither criterion imposes any notion of continuity of W or $\mathbb{X}_f$ with respect to Θ. A constructive design approach exploiting this fact is presented in Section 14.1.1.

13.2.4 Closed-loop Stability

Theorem 13.11 (Main result)**.** *Given system* (19)*, target* Σ_x^o*, and penalty L satisfying Assumptions 13.1, 13.2, 13.3, 13.5, assume the functions* Ψ*,* Ψ_p*,* Ψ_f*,* W *and* $\mathbb{X}_f$ *are designed to satisfy Criteria 13.6, 13.7, 13.9, and 13.10. Furthermore, let* $\mathbb{X}_0 \triangleq \mathbb{X}_0(\Theta^o) \subseteq \mathbb{X}$ *denote the set of initial states, with uncertainty* $\Theta(t_0) = \Theta^o$*, for which* (25) *has a solution. Then under* (23)*,* Σ_x^o *is feasibly asymptotically stabilized from any* $x_0 \in \mathbb{X}_0$*.*

∎

Remark 13.12. *As indicated by Assumption 13.5, the existence of an invariant target set* $\Sigma_x^o(\Theta^o)$*, robust to the full parametric uncertainty* Θ^o*, is not required for Theorem 13.11 to hold. The identifier* $\hat{\Theta}_f$ *must be contained in a sufficiently small neighbourhood of (the worst-case)* θ *such that nontrivial* $\mathbb{X}_f(\hat{\Theta}_f)$ *and* $W(\cdot, \hat{\Theta}_f)$ *exist, for* (25) *to be solvable. While this imposes a minimum performance requirement on* Ψ_f*, it enlarges the domain* $\mathbb{X}_0$ *for which the problem is solvable.*

14. Computation and Performance Issues

14.1 Excitation of the closed-loop trajectories

Contrary to much of the adaptive control literature, including adaptive-MPC approaches such as Mayne & Michalska (1993), the result of Theorem 13.11 does not depend on any auxiliary excitation signal, nor does it require any assumptions regarding the persistency or quality of excitation in the closed-loop behaviour.

Instead, any benefits to the identification which result from injecting excitation into the input signal are predicted by (25c) and (25g), and thereby are automatically accounted for in the posed optimization. In the particular case where $\Psi_p \equiv \Psi_f \equiv \Psi$, then the controller generated by (25) will automatically inject the exact type and amount of excitation which optimizes the cost $J^*(x, \Theta)$; i.e., the closed-loop behaviour (23) could be considered "optimally-exciting". Unlike most *a-priori* excitation signal design methods, excess actuation is not wasted in trying to identify parameters which have little impact on the closed-loop performance (as measured by J^*).

As Ψ_p and Ψ_f deviate from Ψ, the convergence result of Theorem 13.11 remains valid. However, the non-smoothness of $J^*(x, \Theta)$ (with respect to both x and Θ) makes it difficult to quantify the impact of these deviations on the closed-loop behaviour. Qualitatively, small changes to Ψ_p or Ψ_f yielding increasingly conservative identification would generally result in the optimal control solution injecting additional excitation to compensate for the de-sensitized identifier. However, if the changes to Ψ_p or Ψ_f are sufficiently large such that the injection of additional excitation is insufficient to prevent a discontinuous increase in J^*, then it is possible that the optimal solution may suddenly involve *less* excitation than previously, to instead reduce actuation energy. Clearly this behaviour is the result of nonconvexities in the optimal control problem (24), which is inherently a nonconvex problem even in the absence of the adaptive mechanisms proposed here.

14.1.1 A Practical Design Approach for W and $\mathbb{X}_f$

Proposition 14.1. *Let $\{(W^i, \mathbb{X}_f^i)\}$ denote a finitely-indexed collection of terminal function candidates, with indices $i \in \mathcal{I}$, where each pair $(W^i, \mathbb{X}_f^i)$ satisfies Criteria 13.9 and 13.10. Then*

$$W(x, \Theta) \triangleq \min_{i \in \mathcal{I}} \{W^i(x, \Theta)\}, \qquad \mathbb{X}_f(\Theta) \triangleq \bigcup_{i \in \mathcal{I}} \{\mathbb{X}_f^i(\Theta)\} \qquad (27)$$

satisfy Criteria 13.9 and 13.10. ∎

Using Proposition 14.1, it is clear that one approach to constructing $W(\cdot, \cdot)$ and $\mathbb{X}_f(\cdot)$ is to use a collection of pairs of the form

$$\left(W^i(x, \Theta), \mathbb{X}_f^i(\Theta)\right) = \begin{cases} \left(W^i(x), \mathbb{X}_f^i\right) & \Theta \subseteq \Theta^i \\ (+\infty, \varnothing) & \text{otherwise} \end{cases}$$

This collection can be obtained as follows:

1. **Generate a finite collection $\{\Theta^i\}$ of sets covering Θ^o**

   - The elements of the collection can, and should, be overlapping, nested, and ranging in size.

   - Categorize $\{\Theta^i\}$ in a hierarchical (i.e., "tree") structure such that

    i. level 1 (i.e., the top) consists of Θ^o. (Assuming $\Theta^o \in \{\Theta^i\}$ is w.l.o.g., since $W(\cdot, \Theta^o) \equiv +\infty$ and $\mathbb{X}_f(\Theta^o) = \varnothing$ satisfy Criteria 13.9 and 13.10)

    ii. every set in the l'th vertical level is nested inside one or more "parents" on level $l - 1$

    iii. at every level, the "horizontal peers" constitute a cover[12] of Θ^o.

2. **For every set $\Theta^j \in \{\Theta^i\}$, calculate a robust CLF $W^j(\cdot) \equiv W^j(\cdot, \Theta^j)$, and approximate its domain of attraction $\mathbb{X}_f^j \equiv \mathbb{X}_f^j(\Theta^j)$.**

    • Generally, $W^j(\cdot, \Theta^j)$ is selected first, after which $\mathbb{X}_f(\Theta^j)$ is approximated as either a maximal level set of $W^j(\cdot, \Theta^j)$, or as some other controlled-invariant set.

    • Since the elements of $\{\Theta^i\}$ need not be unique, one could actually define multiple $(W^i, \mathbb{X}_f^i)$ pairs associated with the same Θ^j.

    • While not an easy task, this is a very standard robust-control calculation. As such, there is a wealth of tools in the robust control and viability literatures (see, for example Aubin (1991)) to tackle this problem.

3. **Calculate $W(\cdot, \Theta)$ and $\mathbb{X}_f(\Theta)$ online:**

    i. Given Θ, identify all sets that are active: $\mathcal{I}^* = \mathcal{I}^*(\Theta) \triangleq \left\{ j \mid \Theta \subseteq \Theta^j \right\}$. Using the hierarchy, test only immediate children of active parents.

    ii. Given x, search over the active indices to identify $\mathcal{I}_f^* = \mathcal{I}_f^*(x, \mathcal{I}^*) \triangleq \{j \in \mathcal{I}^* \mid x \in \mathbb{X}_f^j\}$. Define $W(x, \Theta) \triangleq \min_{j \in \mathcal{I}_f^*} W^j(x)$ by testing indices in $\mathcal{I}_f^*$, setting $W(x, \Theta) = +\infty$ if $\mathcal{I}_f^* = \varnothing$.

Remark 14.2. *Although the above steps assume that Θ^j is selected before $\mathbb{X}_f^j$, an alternative approach would be to design the candidates $W^j(\cdot)$ on the basis of a collection of parameter values $\hat{\theta}^j$. Briefly, this could be constructed as follows:*

    1. Generate a grid of values $\{\theta^i\}$ distributed across Θ^o.

    2. Design $W^j(\cdot)$ based on a certainty-equivalence model for $\hat{\theta}^j$ (for example, by linearization). Specify $\mathbb{X}_f^j$ (likely as a level set of W^j), and then approximate the maximal neighbourhood Θ^j of $\hat{\theta}^j$ such that Criterion 13.9 holds.

    3. For the same (θ^j, W^j) pair, multiple $(W^j, \mathbb{X}_f^j)$ candidates can be defined corresponding to different Θ^j.

14.2 Robustness Issues

One could argue that if the disturbance model $\mathcal{D}$ in (19) encompasses *all* possible sources of model uncertainty, then the issue of robustness is completely addressed by the min-max formulation of (25). In practice this is not realistic, since it is generally desirable to explicitly consider significant disturbances only, or to exclude $\mathcal{D}$ entirely if Θ represents the dominant uncertainty. The lack of nominal robustness to model error in constrained nonlinear MPC is a well documented problem, as discussed in Grimm et al. (2004). In particular, Grimm et al.

---

[12] specifically, the interiors of all peers must together constitute an open cover

(2003); Marruedo et al. (2002) establish nominal robustness (for "accurate-model", discrete-time MPC) in part by implementing the constraint $x \in \mathbb{X}$ as a succession of strictly nested sets. We present here a modification to this approach, that is relevant to the current adaptive framework.

In addition to ensuring robustness of the controller itself, using methods similar to those mentioned above, it is equally important to ensure that the adaptive mechanism Ψ, including its internal models Ψ_f and Ψ_p, exhibits at least some level of nominal robustness to unmodelled disturbances. By Criterion 13.6.4, the online estimation *must* evolve in a nested fashion and therefore the true θ must never be inadvertently excluded from the estimated uncertainty set. Therefore, just as $\mathbb{Z}$ in (22) defined a best-case bound around which the identifiers in the previous sections could be designed, we present here a modification of (22) which quantifies the type of conservatism required in the identification calculations for the identifiers to possess nominal robustness.

For any $\gamma, \delta \geq 0$, and with $\tau_a \triangleq \tau - a$, we define the following modification of (22):

$$\mathbb{Z}^{\delta,\gamma}(\Theta, x_{[a,b]}, u_{[a,b]}) \triangleq \{\theta \in \Theta \mid B(\dot{x}, \delta + \gamma\tau_a) \cap f(B(x, \gamma\tau_a), u, \theta, \mathcal{D}) \neq \varnothing, \ \forall \tau\}. \qquad (28)$$

Equation (28) provides a conservative outer-approximation of (22) such that $\mathbb{Z} \subseteq \mathbb{Z}^{\delta,\gamma}$. The definition in (28) accounts for two different types of conservatism that can be introduced into the identification calculations. First, the parameter $\delta > 0$ represents a minimum tolerance for the distance between actual derivative information from trajectory $x_{[a,b]}$ and the model (19) when determining if a parameter value can be excluded from the uncertainty set. For situations where the trajectory $x_{[a,b]}$ is itself a *prediction*, as is the case for the internal models Ψ_f and Ψ_p, the parameter $\gamma > 0$ represents increasingly relaxed tolerances applied along the length of the trajectory. Throughout the following we denote $\mathbb{Z}^{\delta} \equiv \mathbb{Z}^{\delta,0}$, with analogous notations for Ψ, Ψ_f, and Ψ_p.

The following technical property of definition (28) is useful towards establishing the desired robustness claim:

Claim 14.3. *For any $a < b < c$, $\gamma \geq 0$, and $\delta \geq \delta' \geq 0$, let $x'_{[a,c]}$ be an arbitrary, continuous perturbation of $x_{[a,b]}$ satisfying*

$$i. \quad \|x'(\tau) - x(\tau)\| \leq \begin{cases} \gamma(\tau - a) & \tau \in [a,b] \\ \gamma(b - a) & \tau \in [b,c] \end{cases}$$

$$ii. \quad \|\dot{x}'(\tau) - \dot{x}(\tau)\| \leq \begin{cases} \delta - \delta' + \gamma(\tau - a) & \tau \in [a,b] \\ \gamma(b - a) & \tau \in [b,c] \end{cases}$$

Then, $\mathbb{Z}^{\delta,\gamma}$ satisfies

$$\mathbb{Z}^{\delta,\gamma}\left(\mathbb{Z}^{\delta'}(\Theta, x'_{[a,b]}, u_{[a,b]}), x'_{[b,c]}, u_{[b,c]}\right) \subseteq \mathbb{Z}^{\delta,\gamma}(\Theta, x_{[a,c]}, u_{[a,c]}). \qquad (29)$$

∎

Based on (28), we are now able to detail sufficient conditions under which the stability claim of Theorem 13.11 holds in the presence of small, unmodelled disturbances. For convenience, the following proposition is restricted to the situation where the only discontinuities in $W(x, \Theta)$ and $\mathbb{X}_f(\Theta)$ are those generated by a switching mechanism (as per Prop. 14.1) between a set of

candidates $\{W^i(x,\Theta), \mathbb{X}_f^i(\Theta)\}$ that are individually *continuous* on $x \in \mathbb{X}_f^i(\Theta)$ (i.e., a strengthening of C13.9.2). With additional complexity, the proposition can be extended to general $\mathcal{LS}$-continuous penalties $W(x,\Theta)$.

Proposition 14.4. *Assume that the following modifications are made to the design in Section 13.2:*

i. *$W(x,\Theta)$ and $\mathbb{X}_f(\Theta)$ are constructed as per Prop. 14.1, but with C13.9.2 strengthened to require the individual $W^i(x,\Theta)$ to be continuous w.r.t $x \in \mathbb{X}_f^i(\Theta)$.*

ii. *For some design parameter $\delta_x > 0$, (26) and (27) are redefined as:*

$$\tilde{L}(\tau, x, u) = \begin{cases} L(x, u) & (x, u) \in \overleftarrow{B}\,(\mathbb{X}, \delta_x \frac{\tau}{T}) \times \mathbb{U} \\ +\infty & \text{otherwise} \end{cases}$$

$$\tilde{W}^i(x, \Theta) = \begin{cases} W^i(x) & x \in \overleftarrow{B}\,(\mathbb{X}_f^i(\Theta), \delta_x) \\ +\infty & \text{otherwise} \end{cases}$$

iii. *The individual sets $\mathbb{X}_f^i$ are specified such that there exists $\delta_f > 0$, for which C13.9.4 holds for every inner approximation $\overleftarrow{B}\,(\mathbb{X}_f^i(\Theta), \delta_x'),\ \delta_x' \in [0, \delta_x]$, where positive invariance is with respect to all flows generated by the differential inclusion $\dot{x} \in B(f(x, k_f^i(x, \Theta), \Theta, \mathcal{D}), \delta_f)$*

iv. *Using design parameters $\delta > \delta' > 0$ and $\gamma > 0$, the identifiers are modified as follows:*

   • *Ψ in (23b) is replaced by $\Psi^{\delta'} \equiv \Psi^{\delta',0}$*

   • *Ψ_p and Ψ_f in (25) are replaced by $\Psi_p^{\delta,\gamma}$ and $\Psi_f^{\delta,\gamma}$, respectively*

   where the new identifiers are assumed to satisfy C13.6, C13.7, and a relation of the form (29).

Then for any compact subset $\bar{\mathbb{X}}_0 \subseteq \mathbb{X}_0(\Theta^o)$, there exists $c^* = c^*(\gamma, \delta_x, \delta_f, \delta, \delta', \bar{\mathbb{X}}_0) > 0$ such that, for all $x_0 \in \bar{\mathbb{X}}_0$ and for all disturbances $\|d_2\| \leq c \leq c^*$, the target Σ_x^o and the actual dynamics

$$\dot{x} = f(x, \kappa_{mpc}(x, \Theta(t)), \theta, d(t)) + d_2(t), \quad x(t_0) = x_0 \tag{30a}$$

$$\Theta(t) = \Psi^{\delta'}(\Theta^o, x_{[t_0,t]}, u_{[t_0,t]}) \tag{30b}$$

are input-to-state stable (ISS); i.e., there exists $\alpha_d \in \mathcal{K}$ such that $x(t)$ asymptotically converges to $B(\Sigma_x^o, \alpha_d(c))$. ∎

14.3 Example Problem

To demonstrate the versatility of our approach, we consider the following nonlinear system:

$$\dot{x}_1 = -x_1 + \left|2\sin(x_1 + \pi\theta_1) + 1.5\theta_2 - x_1 + x_2\right| x_1 + d_1(t)$$

$$\dot{x}_2 = 10\,\theta_{4a}\theta_{4b}x_1\,(u + \theta_3) + d_2(t)$$

The uncertainty $\mathcal{D}$ is given by $|d_1|, |d_2| \leq 0.1$, and Θ^o by $\theta_1, \theta_2, \theta_3 \in [-1,1]$, and $\theta_{4a} \in \{-1,+1\}$, $\theta_{4b} \in [0.5,1]$. The control objective is to achieve regulation of x_1 to the set $x_1 \in [-0.2, 0.2]$, subject to the constraints $\mathbb{X} \triangleq \{|x_1| \leq M_1 \text{ and } |x_2| \leq M_2\}$, $\mathbb{U} \triangleq \{|u| \leq M_u\}$, with $M_1, M_2 \in (0, +\infty]$ and $M_u \in (1, +\infty]$ any given constants. The dynamics exhibit several challenging properties: i) state constraints, ii) nonlinear parameterization of θ_1 and θ_2, iii) potential open-loop instability with finite escape, iv) uncontrollable linearization, v) unknown

sign of control gain, and vi) exogenous disturbances. This system is not stabilizable by any non-adaptive approach (MPC or otherwise), and furthermore fits very few, if any, existing frameworks for adaptive control.

One key property of the dynamics (which is arguably necessary for the regulation objective to be well-posed) is that for any *known* $\theta \in \Theta$ the target is stabilizable and nominally robust. This follows by observing that the surface

$$s \triangleq 2\sin(x_1 + \pi\theta_1) + 1.5\theta_2 - x_1 + x_2 = 0$$

defines a sliding mode for the system, with a robustness margin $|s| \leq 0.5$ for $|x_1| \geq 0.2$. This motivates the design choices:

$$\mathbb{X}_f(\Theta) \triangleq \{x \in \mathbb{X} \mid -M_2 \leq \underline{\Gamma}(x_1, \Theta) \leq x_2 \leq \overline{\Gamma}(x_1, \Theta) \leq M_2\}$$

$$\underline{\Gamma} \triangleq x_1 - 1.5\underline{\theta}_2 - 2\sin(x_1 + \pi\theta_1^{avg}) - 2\pi(\underline{\theta}_1 - \theta_1^{avg}) + 0.5$$

$$\overline{\Gamma} \triangleq x_1 - 1.5\overline{\theta}_2 - 2\sin(x_1 + \pi\theta_1^{avg}) - 2\pi(\overline{\theta}_1 - \theta_1^{avg}) - 0.5$$

where $\overline{\theta}_i, \underline{\theta}_i$ denote upper and lower bounds corresponding to $\Theta \subseteq \Theta^o$, and $\theta^{avg} \triangleq \frac{\overline{\theta} + \underline{\theta}}{2}$. The set $\mathbb{X}_f(\Theta)$ satisfies C13.10 and is nonempty for any Θ such that $\overline{\theta}_2 - \underline{\theta}_2 + \pi(\overline{\theta}_1 - \underline{\theta}_1) \leq 0.5$, that defines minimum thresholds for the performance of Ψ_f and the amount of excitation in solutions to (25).

It can be shown that $|s| \leq 0.5 \overset{\forall \theta \in \Theta^o}{\Longrightarrow} |x_1 - x_2| \leq 4$, and that $\mathbb{X}_f(\Theta)$ is control-invariant using $u \in [-1, 1]$, as long as the sign θ_{4a} is known. This motivates the definitions $\Sigma_u^o \triangleq [-1, 1]$, $\Sigma_1 = [-0.2, 0.2]$, $\Sigma_{12} = [-4, 4]$, and $\Sigma_x^o \triangleq \{x \mid (x_1, x_1 - x_2) \in \Sigma_1 \times \Sigma_{12}\}$, plus the modification of $\mathbb{X}_f(\Theta)$ above to contain the explicit requirement $\Theta_{4a} = \{-1, +1\} \implies \mathbb{X}_f(\Theta) = \emptyset$. Then on $x \in \mathbb{X}_f(\Theta)$, the cost functions $W(x, \Theta) \triangleq \frac{1}{2}\|x_1\|_{\Sigma_1}^2$ and $L(x, u) \triangleq \frac{1}{2}\left(\|x_1\|_{\Sigma_1}^2 + \|x_1 - x_2\|_{\Sigma_{12}}^2 + \|u\|_{\Sigma_u^o}^2\right)$ satisfy all the claims of C13.9, since $W \equiv L \equiv 0$ on $x \in \mathbb{X}_f \cap \Sigma_x^o$, and on $x \in \mathbb{X}_f \setminus \Sigma_x^o$ one has:

$$\dot{W} \leq \|x_1\|_{\Sigma_1}\left(-\tfrac{1}{2}|x_1| + 0.1\right) \leq -\tfrac{1}{2}\|x_1\|_{\Sigma_1}^2 \leq -L(x, u).$$

15. Conclusions

In this chapter we have demonstrated the methodology for adaptive MPC, in which the adverse effects of parameter identification error are explicitly minimized using a robust MPC approach. As a result, it is possible to address both state and input constraints within the adaptive framework. Another key advantage of this approach is that the effects of future parameter estimation can be incorporated into the optimization problem, raising the potential to significantly reduce the conservativeness of the solutions, especially with respect to design of the terminal penalty. While the results presented here are conceptual, in that they are generally intractable to compute due to the underlying min-max feedback-MPC framework, this chapter provides insight into the maximum performance that could be attained by incorporating adaptation into a robust-MPC framework.

16. Proofs for Section 13

16.1 Proof of Theorem 13.11

This proof will follow the so-called "direct method" of establishing stability by directly proving strict decrease of $J^*(x(t), \Theta(t))$, for all $x \notin \Sigma_x^o$. Stability analysis involving $\mathcal{LS}$-continuous Lyapunov functions (for example, (Clarke et al., 1998, Thm4.5.5)) typically involves the proximal subgradient $\partial_p J^*$ (a generalization of ∇J), which is a somewhat ambiguous quantity in the context here given (23b). Instead, this proof exploits an alternative framework involving subderivates (generalized Dini-derivatives), which is equivalent by (Clarke et al., 1998, Prop4.5.3). Together, the following two conditions can be shown sufficient to ensure decrease of J^*, where $\mathcal{F} \triangleq f(x, \kappa_{mpc}(x, \Theta(t)), \Theta(t), \mathcal{D})$

$$(i.) \quad \max_{f \in \mathcal{F}} \overrightarrow{D} J^*(x, \Theta) \triangleq \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{J^*(x + \delta v, \Theta(t + \delta)) - J^*(x, \Theta(t))}{\delta} < 0$$

$$(ii.) \quad \min_{f \in \mathcal{F}} \overleftarrow{D} J^*(x, \Theta) \triangleq \min_{f \in \mathcal{F}} \limsup_{\substack{v \to f \\ \delta \downarrow 0}} \frac{J^*(x - \delta v, \Theta(t - \delta)) - J^*(x, \Theta(t))}{\delta} > 0$$

i.e., J^* is decreasing on both open future and past neighborhoods of t, for all $t \in \mathbb{R}$, where $\overrightarrow{D} J^*, \overleftarrow{D} J^* \in [-\infty, +\infty]$.

To prove condition (i.), let x^p, L^p, W^p, $\hat{\Theta}^p$ correspond to any worst-case minimizing solution of $J^*(x(t), \Theta(t))$, defined on $\tau \in [0, T]$. Additional notations which will be used: $T_\delta \triangleq T + \delta$, $\hat{\Theta}_T^p \triangleq \hat{\Theta}_f(T)$, $\hat{\Theta}_{T_\delta}^p \triangleq \hat{\Theta}_f(T_\delta)$; i.e., both sets represent solutions of the terminal identifier Ψ_f, evaluated along $x_{[0,T]}^p$ and $x_{[0,T_\delta]}^p$, respectively. Likewise, for an arbitrary argument $S \in \{\hat{\Theta}_T^p, \hat{\Theta}_{T_\delta}^p\}$, we define $W_T^p(S) \triangleq W(x^p(T), S)$ and $W_{T_\delta}^p(S) \triangleq W(x^p(T_\delta), S)$.

With the above notations, it can be seen that if the minimizing solution $x_{[0,T]}^p$ were extended to $\tau \in [0, T_\delta]$ by implementing the feedback $u^p(\tau) = k_f(x^p(\tau), \hat{\theta}_T^p)$ on $\tau \in [T, T_\delta]$ (i.e., with $\hat{\theta}_T^p$ fixed), then Criterion C13.9.5 guarantees the inequality

$$\lim_{\delta \downarrow 0} \frac{1}{\delta} \left(\delta L(x_T^p, k_f(x_T^p, \hat{\Theta}_T^p)) + W_{T_\delta}^p(\hat{\Theta}_T^p) - W_T^p(\hat{\Theta}_T^p) \right) \le 0.$$

Using this fact, the relationship (i.) follows from:

$$\max_{f \in \mathcal{F}} \overrightarrow{D} J^*(x, \Theta) = \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[J^*(x + \delta v, \Theta(t + \delta)) - \int_0^T L^p d\tau - W_T^p(\hat{\Theta}_T^p) \right]$$

$$\le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[J^*(x + \delta v, \Theta(t + \delta)) - \int_0^\delta L^p d\tau - \int_\delta^T L^p d\tau - W_T^p(\hat{\Theta}_T^p) \right.$$

$$\left. - \left(\delta L(x_T^p, k_f(x_T^p, \hat{\Theta}_T^p)) + W_{T_\delta}^p(\hat{\Theta}_T^p) - W_T^p(\hat{\Theta}_T^p) \right) \right]$$

$$\le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[J^*(x + \delta v, \Theta(t + \delta)) - \int_\delta^T L^p d\tau - \int_T^{T_\delta} L^p d\tau - W_{T_\delta}^p(\hat{\Theta}_T^p) - \delta L^p|_\delta \right]$$

$$\le \max_{f \in \mathcal{F}} \lim_{\delta \downarrow 0} \frac{1}{\delta} \left[J^*(x^p(\delta), \hat{\Theta}^p(\delta)) - \int_\delta^{T_\delta} L^p d\tau - W_{T_\delta}^p(\hat{\Theta}_{T_\delta}^p) \right] - \delta L^p|_\delta$$

$$\le -L(x, \kappa_{mpc}(x, \Theta))$$

The final inequalities are achieved by recognizing:

- the $\int L^p d\tau + W^p$ term is a (potentially) suboptimal cost on the interval $[\delta, T_\delta]$, starting from the point $(x^p(\delta), \hat{\Theta}_p(\delta))$.

- The relation $\hat{\Theta}^p_{T_\delta} \subseteq \hat{\Theta}^p_T$ holds by Criterion C13.6.4, which implies by Criterion C13.10.2 that $W^p_{T_\delta}(\hat{\Theta}^p_{T_\delta}) \leq W^p_{T_\delta}(\hat{\Theta}^p_T)$

- by C13.7, $\Theta(t+\delta) \triangleq \Psi(\Theta(t), x_{[0,\delta]}, u_{[0,\delta]}) \subseteq \Psi_p(\Theta(t), x_{[0,\delta]}, u_{[0,\delta]})$, along any locus connecting x and $x + \delta v$.

- the $\liminf_v$ applies over all sequences $\{v_k\} \to f$, of which the particular sequence $\{v(\delta_k) = \frac{x^p(\delta_k)-x}{\delta}\}$ is a member.

- there exists an arbitrary perturbation of the sequence $\{v(\delta_k)\}$ satisfying $\Psi_p(\Theta(t), x_{[0,\delta]}) = \hat{\Theta}^p(\delta)$. The $\liminf_v$ includes the limiting cost $J^*(x^p(\delta), \hat{\Theta}^p(\delta))$ of any such perturbation of $\{v(\delta_k)\}$.

- The cost $J^*(x^p(\delta), \hat{\Theta}^p(\delta))$ is optimal on $[\delta, T_\delta]$, and passes through the same point $(x^p(\delta), \hat{\Theta}_p(\delta))$ as the trajectory defining the L^p and W^p expressions. Thus, the bracketed expression is non-positive.

For the purposes of condition (ii.), let x^v denote a solution to the prediction model (25b) for initial condition $x^v(-\delta) = x - \delta v$. Condition (ii.) then follows from:

$$\min_{f \in \mathcal{F}} \overleftarrow{D} J^*(x, \Theta) = \min_{f \in \mathcal{F}} \limsup_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[\int_{-\delta}^{T-\delta} L^v d\tau + W^v_{T-\delta}(\hat{\Theta}^v_{T-\delta}) - J^*(x, \Theta) \right]$$

$$\geq \min_{f \in \mathcal{F}} \limsup_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[\delta L^v|_{-\delta} + \int_0^{T-\delta} L^v d\tau + W^v_{T-\delta}(\hat{\Theta}^v_{T-\delta}) - J^*(x, \Theta) \right.$$
$$\left. + \left(\delta L(x^v_{T-\delta}, k_f(x^v_{T-\delta}, \hat{\Theta}^v_{T-\delta})) + W^v_T(\hat{\Theta}^v_{T-\delta}) - W^v_{T-\delta}(\hat{\Theta}^v_{T-\delta}) \right) \right]$$

$$\geq \min_{f \in \mathcal{F}} \limsup_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[\delta L^v|_{-\delta} + \int_0^T L^v d\tau + W^v_T(\hat{\Theta}^v_{T-\delta}) - J^*(x, \Theta) \right]$$

$$\geq \min_{f \in \mathcal{F}} \lim_{\delta \downarrow 0} \frac{1}{\delta} \left[\delta L^p|_{-\delta} + \int_0^T L^p d\tau + W^p_T(\hat{\Theta}^p_T) - J^*(x, \Theta) \right]$$

$$\geq L(x, \kappa_{mpc}(x, \Theta))$$

The above derivation made use of the fact that the reverse subderivate $\overleftarrow{D} W$ satisfies

$$\min_{f \in \mathcal{F}} \limsup_{\substack{v \to f \\ \delta \downarrow 0}} \left(-L(x-\delta v, k_f(x-\delta v, \Theta)) + \left(\frac{W(x-\delta v, \Theta) - W(x, \Theta)}{\delta} \right) \right) \geq 0$$

which follows from a combination of C13.9.5 and the $\mathcal{LS}$-continuity of W.

Using the above inequalities for $\overleftarrow{D} J^*(x, \Theta)$ and $\overrightarrow{D} J^*(x, \Theta)$ together with Assumption 13.3, it follows that $J^*(t)$ is strictly decreasing on $x \notin \Sigma^o_x$ and non-increasing on $x \in \Sigma^o_x$. It follows that $\lim_{t \to \infty}(x, \Theta)$ must converge to an invariant subset of $\Sigma^0_x \times \overline{\text{cov}}\{\Theta^o\}$. Assumption 13.1 guarantees that such an invariant subset exists, since it implies $\exists \delta^* > 0$ such $\Sigma_x(B(\theta^*, \delta^*)) \neq \emptyset$, with θ^* the actual unknown parameter in (19). Continued solvability of (25) as $(x(t), \Theta(t))$ evolve follows by: 1) $x(\tau) \notin \mathbb{X}_0(\Theta(\tau)) \Rightarrow J^*(\tau) = +\infty$, and 2) if $x(t) \in \mathbb{X}_0(\Theta(t))$ and $x(t') \notin \mathbb{X}_0(\Theta(t'))$, then $(t' - t) \downarrow 0$ contradicts either condition (i.) at time t, or (ii.) at time t'.

16.2 Proof of Proposition 14.1

The fact that C13.10 holds is a direct property of the union and min operations for the closed sets $\mathbb{X}_f^i$, and the fact that the Θ-dependence of individual $(W^i, \mathbb{X}_f^i)$ satisfies C13.10. For the purposes of C13.9, the Θ argument is a constant, and is omitted from notation. Properties C13.9.1 and C13.9.2 follow directly by (27), the closure of $\mathbb{X}_f^i$, and (2). Define

$$\mathcal{I}_f(x) \,\&\, = \{i \in \mathcal{I} \mid x \in \mathbb{X}_f^i \text{ and } W(x) = W^i(x)\}$$

Denoting $\mathcal{F}^i \triangleq f(x, k_f^i(x), \Theta, \mathcal{D})$, the following inequality holds for every $i \in \mathcal{I}_f(x)$:

$$\max_{f^i \in \mathcal{F}^i} \liminf_{\substack{v \to f^i \\ \delta \downarrow 0}} \frac{W(x + \delta v) - W(x)}{\delta} \quad \leq \quad \max_{f^i \in \mathcal{F}^i} \liminf_{\substack{v \to f^i \\ \delta \downarrow 0}} \frac{W^i(x + \delta v) - W(x)}{\delta} \quad \leq \quad -L(x, k_f^i(x))$$

It then follows that $u = k_f(x) \triangleq k_f^{i(x)}(x)$ satisfies C13.9.5 for any arbitrary selection rule $i(x) \in \mathcal{I}_f(x)$ (from which C13.9.3 is obvious). Condition C13.9.4 follows from continuity of the $x(\cdot)$ flows, and observing that by (26), C13.9.5 would be violated at any point of departure from $\mathbb{X}_f$.

16.3 Proof of Claim 14.3

By contradiction, let θ^* be a value contained in the left-hand side of (29), but not in the right-hand side. Then by (28), there exists $\tau \in [a, c]$ (i.e., $\tau_a \equiv (\tau - a) \in [0,\ c - a]$) such that

$$f(B(x, \gamma\tau_a), u, \theta^*, \mathcal{D}) \cap B(\dot{x}, \delta + \gamma\tau_a) = \varnothing \tag{31}$$

Using the bounds indicated in the claim, the following inclusions hold when $\tau \in [a, b]$:

$$f(x', u, \theta^*, \mathcal{D}) \subseteq f(B(x, \gamma\tau_a), u, \theta^*, \mathcal{D}) \tag{32a}$$

$$B(\dot{x}', \delta') \subseteq B(\dot{x}, \delta + \gamma\tau_a) \tag{32b}$$

Combining (32) and (31) yields

$$f(x', u, \theta^*, \mathcal{D}) \cap B(\dot{x}', \delta') = \varnothing \quad \implies \quad \theta^* \notin \mathbb{Z}^{\delta'}(\Theta, x'_{[a,\tau]}, u_{[a,\tau]}) \tag{33}$$

which violates the initial assumption that θ^* is in the LHS of (29). Meanwhile, for $\tau \in [b, c]$ the inclusions

$$f(B(x', \gamma\tau_b), u, \theta^*, \mathcal{D}) \subseteq f(B(x, \gamma\tau_a), u, \theta^*, \mathcal{D}) \tag{34a}$$

$$B(\dot{x}', \delta + \gamma\tau_b) \subseteq B(\dot{x}, \delta + \gamma\tau_a) \tag{34b}$$

yield the same contradictory conclusion:

$$f(B(x', \gamma\tau_b), u, \theta^*, \mathcal{D}) \cap B(\dot{x}', \delta + \gamma\tau_b) = \varnothing \tag{35a}$$

$$\implies \theta^* \notin \mathbb{Z}^{\delta, \gamma}\left(\mathbb{Z}^{\delta'}(\Theta, x'_{[a,b]}, u_{[a,b]}), x'_{[b,\tau]}, u_{[b,\tau]}\right) \tag{35b}$$

It therefore follows that the containment indicated in (29) necessarily holds.

16.4 Proof of Proposition 14.4

It can be shown that Assumption 13.3, together with the compactness of Σ_x, is sufficient for an analogue of Claim **??** to hold (i.e., with J_∞^* interpreted in a $\min - \max$ sense). In other words, the cost $J^*(x, \Theta)$ satisfies

$$\alpha_l(\|x\|_{\Sigma_x^o}, \Theta) \le J^*(x, \Theta) \le \alpha_h(\|x\|_{\Sigma_x^o}, \Theta)$$

for some functions α_l, α_h which are class-$\mathcal{K}_\infty$ w.r.t. x, and whose parameterization in Θ satisfies $\alpha_i(x, \Theta_1) \le \alpha_i(x, \Theta_2), \Theta_1 \subseteq \Theta_2$. We then define the compact set $\bar{\mathbb{X}}_0^\uparrow \triangleq \{x \mid \min_{\Theta \in \overline{\text{cov}}\{\Theta^o\}} J^*(x, \Theta) < \max_{x_0 \in \bar{\mathbb{X}}_0} \alpha_h(\|x_0\|_{\Sigma_x^o}, \Theta^0)\}$.

By a simple extension of (Khalil, 2002, Thm4.19), the ISS property follows if it can be shown that there exists $\alpha_c \in \mathcal{K}$ such that $J^*(x, \Theta)$ satisfies

$$x \in \bar{\mathbb{X}}_0^\uparrow \setminus B(\Sigma_x^o, \alpha_c(c)) \Rightarrow \begin{cases} \max_{f \in \mathcal{F}_c} \overrightarrow{D} J^*(x, \Theta) < 0 \\ \min_{f \in \mathcal{F}_c} \overleftarrow{D} J^*(x, \Theta) > 0 \end{cases} \tag{36}$$

where $\mathcal{F}_c \triangleq B(f(x, \kappa_{mpc}(x, \Theta(t)), \Theta(t), \mathcal{D}), c)$. To see this, it is clear that J decreases until $x(t)$ enters $B(\Sigma_x^o, \alpha_c(c))$. While this set is not necessarily invariant, it is contained within an invariant, compact levelset $\Omega(c, \Theta) \triangleq \{x \mid J^*(x, \Theta) \le \alpha_h(\alpha_c(c), \Theta)\}$. By C13.6.4, the evolution of $\Theta(t)$ in (30b) must approach some constant interior bound Θ^∞, and thus $\lim_{t \to \infty} x(t) \in \Omega(c, \Theta^\infty)$. Defining $\alpha_d(c) \triangleq \max_{x \in \Omega(c, \Theta^\infty)} \|x\|_{\Sigma_x^o}$ completes the Proposition, if c^* is sufficiently small such that $B(\Sigma_x^o, \alpha_d(c^*)) \subseteq \bar{\mathbb{X}}_0^\uparrow$.

Next, we only prove decrease in the forward direction, since the reverse direction follows analogously, as it did in the proof of Theorem 13.11. Using similar procedure and notation as the Thm 13.11 proof, $x_{[0,T]}^p$ denotes any worst-case prediction at (t, x, Θ), extended to $[T, T_\delta]$ via k_f, that is assumed to satisfy the specifications of Proposition 14.4. Following the proof of Theorem 13.11,

$$\max_{f \in \mathcal{F}_{c^*}} \overrightarrow{D} J^*(x, \Theta) \le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \Big[J^*(x + \delta v, \Theta(t + \delta)) - \int_\delta^{T_\delta} L^p d\tau - W_{T_\delta}^p(\hat{\Theta}_T^p) \Big] - L^p|_\delta$$

$$\le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \Big[J^*(x + \delta v, \Theta(t + \delta)) - \int_\delta^{T_\delta} L^v d\tau - W_{T_\delta}^v(\hat{\Theta}_{T_\delta}^v) \Big] - L^p|_\delta$$

$$+ \frac{1}{\delta} \Big[\int_\delta^{T_\delta} L^v d\tau + W_{T_\delta}^v(\hat{\Theta}_{T_\delta}^v) - \int_\delta^{T_\delta} L^p d\tau - W_{T_\delta}^p(\hat{\Theta}_T^p) \Big] \tag{37}$$

where L^v, W^v denote costs associated with a trajectory $x_{[0,T_\delta]}^v$ satisfying the following:

- initial conditions $x^v(0) = x, \Theta^v(0) = \Theta$.

- generated by the same worst-case $\hat{\theta}$ and $d(\cdot)$ as $x_{[0,T_\delta]}^p$

- dynamics of form (30) on $\tau \in [0, \delta]$, and of form (25b),(25c) on $\tau \in [\delta, T_\delta]$, with the trajectory passing through $x^v(\delta) = x + \delta v, \Theta_p^v(\delta) = \Theta(t + \delta)$.

- the $\min_\kappa$ in (25) is constrained such that $\kappa^v(\tau, x^v, \Theta^v) = \kappa^p(\tau, x^p, \Theta^p)$; i.e., $u_{[0,T_\delta]}^v \equiv u_{[0,T_\delta]}^p \equiv u_{[0,T_\delta]}$.

Let K_f denote a Lipschitz constant of (19) with respect to x, over the compact domain $\bar{\mathbb{X}}_0^\uparrow \times \Theta^o \times \mathcal{D}$. Then, using the comparison lemma (Khalil, 2002, Lem3.4) one can derive the bounds

$$\tau \in [0, \delta]: \quad \begin{cases} \|x^v - x^p\| \le \frac{c}{K_f}(e^{K_f \tau} - 1) \\ \|\dot{x}^v - \dot{x}^p\| \le c\, e^{K_f \tau} \end{cases} \tag{38a}$$

$$\tau \in [\delta, T_\delta]: \quad \begin{cases} \|x^v - x^p\| \le \frac{c}{K_f}(e^{K_f \delta} - 1)\, e^{K_f(\tau - \delta)} \\ \|\dot{x}^v - \dot{x}^p\| \le c\,(e^{K_f \delta} - 1)\, e^{K_f(\tau - \delta)} \end{cases} \tag{38b}$$

As $\delta \downarrow 0$, the above inequalities satisfy the conditions of Claim 14.3 as long as $c^* < \min\{\gamma, (\delta - \delta'), \gamma e^{K_f T}, \frac{\gamma}{K_f} e^{K_f T}\}$, thus yielding

$$\hat{\Theta}_f^v = \Psi_f^{\delta,\gamma}(\Psi^{\delta'}(\Theta, x_{[0,\delta]}^v, u_{[0,\delta]}), x_{[\delta,T_\delta]}^v, u_{[\delta,T_\delta]}) \subseteq \Psi_f^{\delta,\gamma}(\Theta, x_{[0,T_\delta]}^p, u_{[0,T_\delta]}) = \hat{\Theta}_f^p$$

as well as the analogue $\hat{\Theta}_p^v(\tau) \subseteq \hat{\Theta}_p^p(\tau)$, $\forall \tau \in [0, T_\delta]$.

Since $x_{[0,T]}^p$ is a feasible solution of the original problem from (t, x, Θ) with $\tau \in [0, T]$, it follows for the new problem posed at time $t + \delta$ that x^v is feasible with respect to the appropriate inner approximations of $\mathbb{X}$ and $\mathbb{X}_f^{i*}(\hat{\Theta}_T^p) \subseteq \mathbb{X}_f(\hat{\Theta}_{T_\delta}^v)$ (where i^* denotes an active terminal set for x_f^p) if

$$\|x^v - x^p\| \le \begin{cases} \delta \frac{\delta_x}{T} & \tau \in [\delta, T] \\ \delta\,\delta_f & \tau \in [T, T_\delta] \end{cases}$$

which holds by (38) as long as $c^* < \min\{\delta_f, \frac{\delta_x}{T}\} e^{-K_f T}$. Using arguments from the proof Theorem 13.11, the first term in (37) can be eliminated, leaving:

$$\begin{aligned}
\max_{f \in \mathcal{F}_c} \overrightarrow{D} J^*(x, \Theta) &\le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left[\int_\delta^{T_\delta} L^v d\tau + W_{T_\delta}^v(\hat{\Theta}_{T_\delta}^v) - \int_\delta^{T_\delta} L^p d\tau - W_{T_\delta}^p(\hat{\Theta}_T^p) \right] - L^p|_\delta \\
&\le \max_{f \in \mathcal{F}} \liminf_{\substack{v \to f \\ \delta \downarrow 0}} \frac{1}{\delta} \left(\int_\delta^{T_\delta} K_L \|x^v - x^p\| d\tau + K_W \|x^v(T) - x^p(T)\| - L^p|_\delta \right) \\
&\le \lim_{\delta \downarrow 0} \left(\frac{c(e^{K_f \delta} - 1)}{K_f \delta} \left[K_W + T K_L \right] e^{K_f T} - L^p|_\delta \right) \\
&\le -L(x, k_{MPC}(x, \Theta)) + c(K_W + T K_L) e^{K_f T} \\
&< 0 \qquad \forall x \in \bar{\mathbb{X}}_0^\uparrow \setminus B(\Sigma_x^o, \alpha_c(c))
\end{aligned}$$

with $\alpha_c \in \mathcal{K}$ given by

$$\alpha_c(c) \triangleq \underline{\gamma}_L^{-1} \left(c\,(K_W + T K_L)\, e^{K_f T} \right)$$

where K_W is a Lipschitz constant of $W^{i*}(x, \Theta)$ over the compact domain $\bar{\mathbb{X}}_0^\uparrow \cap \mathbb{X}_f^{i*}(\Theta)$, maximal over all $\Theta \in \overline{\text{cov}}\{\Theta^o\}$. Likewise, K_L is a Lipschitz constant of $L(x, u)$ with respect to x, maximal over $u \in \mathbb{U}$.

This proves the forward case in (36), with the reverse case following similarly. As argued previously, this is sufficient to yield the ISS property of (30) with respect to $\|d_2\| \le c \le c^*$, which completes the proof.

17. References

Adetola, V. & Guay, M. (2004). Adaptive receding horizon control of nonlinear systems, *Proc. IFAC Symposium on Nonlinear Control Systems*, Stuttgart, Germany, pp. 1055–1060.

Aubin, J. (1991). *Viability Theory*, Systems & Control: Foundations & Applications, Birkhäuser, Boston.

Bellman, R. (1952). The theory of dynamic programming, *Proc. National Academy of Science,*, number 38, USA.

Bellman, R. (1957). *Dynamic Programming*, Princeton Press.

Bertsekas, D. (1995). *Dynamic Programming and Optimal Control*, Vol. I, Athena Scientific, Belmont, MA.

Brogliato, B. & Neto, A. T. (1995). Practical stabilization of a class of nonlinear systems with partially known uncertainties, *Automatica* **31**(1): 145 – 150.

Bryson, A. & Ho, Y. (1969). *Applied Optimal Control*, Ginn and Co., Waltham, MA.

Cannon, M. & Kouvaritakis, B. (2005). Optimizing prediction dynamics for robust MPC, **50**(11): 1892–1897.

Chen, H. & Allgöwer, F. (1998a). A computationally attractive nonlinear predictive control scheme with guaranteed stability for stable systems, *Journal of Process Control* **8**(5-6): 475–485.

Chen, H. & Allgöwer, F. (1998b). A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability, *Automatica* **34**(10): 1205–1217.

Chen, H., Scherer, C. & Allgöwer (1997). A game theoretic approach to nonlinear robust receding horizon control of constrained systems, *Proc. American Control Conference*.

Clarke, F., Ledyaev, Y., Stern, R. & Wolenski, P. (1998). *Nonsmooth Analysis and Control Theory*, Grad. Texts in Math. 178, Springer-Verlag, New York.

Corless, M. J. & Leitmann, G. (1981). Continuous state feedback guaranteeing uniform ultimate boundedness for uncertain dynamic systems., *IEEE Trans. Automat. Contr.* **AC-26**(5): 1139 – 1144.

Coron, J. & Rosier, L. (1994). A relation between continuous time-varying and discontinuous feedback stabilization, *Journal of Mathematical Systems, Estimation, and Control* **4**(1): 67–84.

Cutler, C. & Ramaker, B. (1980). Dynamic matrix control - a computer control algorithm, *Proceedings Joint Automatic Control Conference*, San Francisco, CA.

De Nicolao, G., Magni, L. & Scattolini, R. (1996). On the robustness of receding horizon control with terminal constraints, *IEEE Trans. Automat. Contr.* **41**: 454–453.

Findeisen, R., Imsland, L., Allgöwer, F. & Foss, B. (2003). Towards a sampled-data theory for nonlinear model predictive control, *in* C. Kang, M. Xiao & W. Borges (eds), *New Trends in Nonlinear Dynamics and Control, and their Applications*, Vol. 295, Springer-Verlag, New York, pp. 295–313.

Freeman, R. & Kokotović, P. (1996a). Inverse optimality in robust stabilization, *SIAM Journal of Control and Optimization* **34**: 1365–1391.

Freeman, R. & Kokotović, P. (1996b). *Robust Nonlinear Control Design*, Birkh auser.

Grimm, G., Messina, M., Tuna, S. & Teel, A. (2003). Nominally robust model predictive control with state constraints, *Proc. IEEE Conf. on Decision and Control*, pp. 1413–1418.

Grimm, G., Messina, M., Tuna, S. & Teel, A. (2004). Examples when model predictive control is non-robust, *Automatica* **40**(10): 1729–1738.

Grimm, G., Messina, M., Tuna, S. & Teel, A. (2005). Model predictive control: for want of a local control lyapunov function, all is not lost, *IEEE Trans. Automat. Contr.* **50**(5): 617–628.

Hermes, H. (1967). Discontinuous vector fields and feedback control, *in* J. Hale & J. LaSalle (eds), *Differential Equations and Dynamical Systems*, Academic Press, New York, pp. 155–166.

Hestenes, M. (1966). *Calculus of Variations and Optimal Control*, John Wiley & Sons, New York.

Jadbabaie, A., Yu, J. & Hauser, J. (2001). Unconstrained receding-horizon control of nonlinear systems, *IEEE Trans. Automat. Contr.* **46**(5): 776 – 783.

Kalman, R. (1960). Contributions to the theory of optimal control, *Bol. Soc. Mat. Mexicana* **5**: 102–119.

Kalman, R. (1963). Mathematical description of linear dynamical systems, *SIAM J. Control* **1**: 152–192.

Keerthi, S. S. & Gilbert, E. G. (1988). Optimal, infinite horizon feedback laws for a general class of constrained discrete time systems: Stability and moving-horizon approximations, *Journal of Optimization Theory and Applications* **57**: 265–293.

Khalil, H. (2002). *Nonlinear Systems*, 3rd edn, Prentice Hall, Englewood Cliffs, N.J.

Kim, J.-K. & Han, M.-C. (2004). Adaptive robust optimal predictive control of robot manipulators, *IECON Proceedings (Industrial Electronics Conference)* **3**: 2819 – 2824.

Kothare, M., Balakrishnan, V. & Morari, M. (1996). Robust constrained model predictive control using linear matrix inequalities, *Automatica* **32**(10): 1361–1379.

Kouvaritakis, B., Rossiter, J. & Schuurmans, J. (2000). Efficient robust predictive control, *IEEE Trans. Automat. Contr.* **45**(8): 1545 – 1549.

Langson, W., Chryssochoos, I., Raković, S. & Mayne, D. (2004). Robust model predictive control using tubes, *Automatica* **40**(1): 125 – 133.

Lee, E. & Markus, L. (1967). *Foundations of Optimal Control Theory*, Wiley.

Lee, J. & Yu, Z. (1997). Worst-case formulations of model predictive control for systems with bounded parameters, *Automatica* **33**(5): 763–781.

Magni, L., De Nicolao, G., Scattolini, R. & Allgöwer, F. (2003). Robust model predictive control for nonlinear discrete-time systems, *International Journal of Robust and Nonlinear Control* **13**(3-4): 229–246.

Magni, L., Nijmeijer, H. & van der Schaft, A. (2001). Receding-horizon approach to the nonlinear h_∞ control problem, *Automatica* **37**(3): 429 – 435.

Magni, L. & Sepulchre, R. (1997). Stability margins of nonlinear receding-horizon control via inverse optimality, *Systems and Control Letters* **32**: 241–245.

Marruedo, D., Alamo, T. & Camacho, E. (2002). Input-to-state stable MPC for constrained discrete-time nonlinear systems with bounded additive uncertainties, *Proc. IEEE Conf. on Decision and Control*, pp. 4619–4624.

Mayne, D. (1995). Optimization in model based control, *Proc. IFAC symposium on dynamics and control, chemical reactors and batch processes (DYCORD)*, Oxford: Elsevier Science., pp. 229–242. plenary address.

Mayne, D. Q. & Michalska, H. (1990). Receding horizon control of non-linear systems, *IEEE Trans. Automat. Contr.* **35**(5): 814–824.

Mayne, D. Q. & Michalska, H. (1993). Adaptive receding horizon control for constrained nonlinear systems, *Proc. IEEE Conf. on Decision and Control*, pp. 1286–1291.

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

Michalska, H. & Mayne, D. (1993). Robust receding horizon control of constrained nonlinear systems, *IEEE Trans. Automat. Contr.* **38**(11): 1623 – 1633.

Pontryagin, L. (1961). Optimal regulation processes, *Amer. Math. Society Trans., Series 2* **18**: 321–339.

Primbs, J. (1999). *Nonlinear Optimal Control: A Receding Horizon Approach*, PhD thesis, California Institute of Technology, Pasadena, California.

Primbs, J., Nevistic, V. & Doyle, J. (2000). A receding horizon generalization of pointwise min-norm controllers, *IEEE Trans. Automat. Contr.* **45**(5): 898–909.

Raković, S. & Mayne, D. (2005). Robust time optimal obstacle avoidance problem for constrained discrete time systems, *Proc. IEEE Conf. on Decision and Control*.

Ramirez, D., Alamo, T. & Camacho, E. (2002). Efficient implementation of constrained min-max model predictive control with bounded uncertainties, *Proc. IEEE Conf. on Decision and Control*, pp. 3168–3173.

Richalet, J., Rault, A., Testud, J. & Papon, J. (1976). Algorithmic control of industrial processes, *Proc. IFAC symposium on identification and system parameter estimation*, pp. 1119–1167.

Richalet, J., Rault, A., Testud, J. & Papon, J. (1978). Model predictive heuristic control: Applications to industrial processes, *Automatica* **14**: 413–428.

Sage, A. P. & White, C. C. (1977). *Optimum Systems Control*, 2nd edn, Prentice-Hall.

Scokaert, P. & Mayne, D. (1998). Min-max feedback model predictive control for constrained linear systems, *IEEE Trans. Automat. Contr.* **43**(8): 1136–1142.

Sepulchre, R., Jankovic, J. & Kokotovic, P. (1997). *Constructive Nonlinear Control*, Springer, New York.

Sontag, E. (1989). A "universal" construction of artstein's theorem on nonlinear stabilization, *Systems and Control Letters* **13**: 117–123.

Sontag, E. D. (1983). Lyapunov-like characterization of asymptotic controllability., *SIAM Journal on Control and Optimization* **21**(3): 462 – 471.

Tang, Y. (1996). Simple robust adaptive control for a class of non-linear systems: an adaptive signal synthesis approach, *International Journal of Adaptive Control and Signal Processing* **10**(4-5): 481 – 488.

Tuna, S., Sanfelice, R., Messina, M. & Teel, A. (2005). Hybrid MPC: Open-minded but not easily swayed, *International Workshop on Assessment and Future Directions of Nonlinear Model Predictive Control*, Freudenstadt-Lauterbad, Germany, pp. 169–180.

Model Predictive Control

Edited by Tao Zheng

Frontiers of Model Predictive Control Robust Model Predictive Control Nonlinear Model Predictive Control Excellent Applications Guide for Researchers and Engineers Recent Achievements of Authors over the World Theory with Practical Examples Kinds of Algorithms for Choice

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

INTECH
open science | open minds