



Survey Paper

Constrained model predictive control: Stability and optimality[☆]D. Q. Mayne^{a,*}, J. B. Rawlings^b, C. V. Rao^b, P. O. M. Sokaert^c^aDepartment of Electrical and Electronic Engineering, Imperial College of Science, Technology and Medicine, London SW7 2BT, UK^bDepartment of Chemical Engineering, University of Wisconsin, Madison, USA^cCentre National d'Etudes des Telecommunications, France Telecom, France

Received 29 May 1998; revised 6 September 1999; received in final form 8 November 1999

Abstract

Model predictive control is a form of control in which the current control action is obtained by solving, at *each* sampling instant, a finite horizon open-loop optimal control problem, using the current state of the plant as the initial state; the optimization yields an optimal control sequence and the first control in this sequence is applied to the plant. An important advantage of this type of control is its ability to cope with hard constraints on controls and states. It has, therefore, been widely applied in petro-chemical and related industries where satisfaction of constraints is particularly important because efficiency demands operating points on or close to the boundary of the set of admissible states and controls. In this review, we focus on model predictive control of constrained systems, both linear and nonlinear and discuss only briefly model predictive control of unconstrained nonlinear and/or time-varying systems. We concentrate our attention on research dealing with stability and optimality; in these areas the subject has developed, in our opinion, to a stage where it has achieved sufficient maturity to warrant the active interest of researchers in nonlinear control. We distill from an extensive literature essential principles that ensure stability and use these to present a concise characterization of most of the model predictive controllers that have been proposed in the literature. In some cases the finite horizon optimal control problem solved on-line is exactly equivalent to the same problem with an infinite horizon; in other cases it is equivalent to a modified infinite horizon optimal control problem. In both situations, known advantages of infinite horizon optimal control accrue. © 2000 Elsevier Science Ltd. All rights reserved.

Keywords: Model predictive control; Stability; Optimality; Robustness

1. Introduction*1.1. What is model predictive control?*

Model predictive control (MPC) or receding horizon control (RHC) is a form of control in which the current control action is obtained by solving *on-line*, at *each* sampling instant, a finite horizon open-loop optimal control problem, using the current state of the plant as the initial state; the optimization yields an optimal control sequence and the first control in this sequence is applied

to the plant. This is its main difference from conventional control which uses a pre-computed control law. With the clarity gained by hindsight, it can be recognized that the *raison d'être* for model predictive control is its ability to handle control problems where off-line computation of a control law is difficult or impossible although other features, such as its capability for controlling multivariable plants, were initially deemed more important. Nearly every application imposes constraints; actuators are naturally limited in the force (or equivalent) they can apply, safety limits states such as temperature, pressure and velocity and efficiency often dictates steady-state operation close to the boundary of the set of permissible states. The prevalence of hard constraints is accompanied by a dearth of control methods for handling them, despite a continuous demand from industry that has had, in their absence, to resort often to ad hoc methods. Model predictive control is one of few suitable methods, and this fact makes it an important tool for the control engineer, particularly in the process industries where plants being

[☆]This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Editor M. Morari. Research supported by National Science Foundation under grant ECS-93-12922 and CTS-9708497.

* Corresponding author. Tel.: +44-171-594-6246; fax: +44-171-594-6282.

E-mail address: d.mayne@ic.ac.uk (D. Q. Mayne)

controlled are sufficiently ‘slow’ to permit its implementation. Other examples where model predictive control may be advantageously employed include unconstrained nonlinear plants, for which off-line computation of a control law usually requires the plant dynamics to possess a special structure, and time-varying plants.

MPC is *not* a new method of control design. It essentially solves standard optimal control problems (except that, in MPC, the optimal control problem is required to have a finite horizon in contrast to the infinite horizon usually employed in H_2 and H_∞ linear optimal control). Where it differs from other controllers is that it solves the optimal control problem *on-line* for the current state of the plant, rather than determining *off-line* a feedback policy (that provides the optimal control for all states). The on-line solution is obtained by solving an open-loop optimal control problem in which the initial state is the current state of the system being controlled; this is a mathematical programming problem. Determining the feedback solution, on the other hand, requires solution of the Hamilton–Jacobi–Bellman (Dynamic Programming) differential or difference equation, a vastly more difficult task (except in those cases, such as H_2 and H_∞ linear optimal control, where the value function can be finitely parameterized). From this point of view, MPC differs from other control methods merely in its implementation. The requirement that the open-loop optimal control problem be solvable in a reasonable time (compared with plant dynamics) necessitates, however, the use of a finite horizon and this raises interesting problems.

1.2. The open-loop optimal control problem

The system to be controlled is usually described, or approximated, by an ordinary differential equation but, since the control is normally piecewise constant, is usually modelled, in the MPC literature, by a difference equation

$$x(k + 1) = f(x(k), u(k)), \tag{1.1}$$

$$y(k) = h(x(k)), \tag{1.2}$$

where $f(\cdot)$ is implicitly defined by the originating differential equation that has an equilibrium point at the origin ($f(0, 0) = 0$). In the sequel, we often use x^+ to denote the successor state $f(x, u)$ and write (1.1) and (1.2) in the abbreviated form $x^+ = f(x, u)$, $y = h(x)$. We employ $u(\cdot)$ or \mathbf{u} to denote a control sequence and $x^u(\cdot; (x, j))$ to denote the state trajectory (sequence) resulting from an initial state x at time j and a control sequence \mathbf{u} . The control and state sequences must satisfy

$$u(k) \in \mathbb{U}, \tag{1.3}$$

$$x(k) \in \mathbb{X}, \tag{1.4}$$

where, usually, \mathbb{U} is a convex, compact subset of \mathbb{R}^m and \mathbb{X} a convex, closed subset of \mathbb{R}^n , each set containing the origin in its interior. The control objective is usually to steer the state to the origin or to an equilibrium state x_r for which the output $y_r = h(x_r) = r$ where r is the constant reference. A suitable change of coordinates reduces the second problem to the first which, therefore, we consider in the sequel. For event (x, k) (i.e. for state x at time k), the cost is defined by

$$V(x, k, \mathbf{u}) = \sum_{i=k}^{k+N-1} \ell(x(i), u(i)) + F(x(k + N)), \tag{1.5}$$

where $\mathbf{u} = \{u(k), u(k + 1), \dots, u(k + N - 1)\}$ and $x(i) := x^u(i; (x, k))$. We assume the stage cost $\ell(x, u) \geq c(\|(x, u)\|)^2$ to avoid technicalities (this assumption may be replaced by $\ell(x, u) \geq c(\|z, u\|)^2$ where $z = h(x)$ and (f, h) is detectable) and that $\ell(0, 0) = 0$. The terminal time $k + N$ increases with time k and is often referred to as a *receding horizon*. A terminal constraint

$$x(k + N) \in X_f \subset \mathbb{X} \tag{1.6}$$

is sometimes imposed. At event (x, k) , the optimal control problem $\mathcal{P}(x, k)$ of minimizing $V(x, k, \mathbf{u})$ subject to the control, state and terminal constraints is solved, yielding the optimizing control sequence

$$\mathbf{u}^0(x, k) = \{u^0(k; (x, k)), u^0(k + 1; (x, k)), \dots, u^0(k + N - 1; (x, k))\} \tag{1.7}$$

and the value function

$$V^0(x, k) = V(x, k, \mathbf{u}^0(x, k)) \tag{1.8}$$

(the argument (x, k) denotes the initial state is x at time k). The first control $u^0(k; (x, k))$ in the optimal sequence $\mathbf{u}^0(x, k)$ is applied to the system (at time k). This defines an implicit model predictive control law $\kappa(x, k) := u^0(k; (x, k))$.

Since $f(\cdot)$ and $\ell(\cdot)$ are time invariant, the problems $\mathcal{P}(x, k)$ are time invariant in the sense that $V^0(x, k) = V^0(x, 0)$ and $\kappa(x, k) = \kappa(x, 0)$ for all k , so that it suffices, at each event (x, k) to solve $\mathcal{P}_N(x) := \mathcal{P}(x, 0)$, i.e. to regard current time as zero. Problem $\mathcal{P}_N(x)$ (suffix N is used to remind us that the optimization is over a finite horizon N) is therefore defined by

$$\mathcal{P}_N(x): V_N^0(x) = \min_{\mathbf{u}} \{V_N(x, \mathbf{u}) \mid \mathbf{u} \in \mathcal{U}_N(x)\}, \tag{1.9}$$

where, now,

$$V_N(x, \mathbf{u}) := \sum_{i=0}^{N-1} \ell(x(i), u(i)) + F(x(N)) \tag{1.10}$$

$\mathbf{u} = \{u(0), u(1), \dots, u(N - 1)\}$, $x(i) = x^u(i; (x, 0))$ and $\mathcal{U}_N(x)$ is the set of (feasible) control sequences satisfying the

control, state and terminal constraints. Because N is finite, the minimum exists if $f(\cdot)$, $\ell(\cdot)$ and $F(\cdot)$ are continuous, \cup compact, and \times and X_f closed. At event (x, k) , the problem $\mathcal{P}_N(x)$ is solved yielding the optimizing control sequence

$$\mathbf{u}^0(x) = \{u^0(0; x), u^0(1; x), \dots, u^0(N-1; x)\} \quad (1.11)$$

the optimal state trajectory

$$\mathbf{x}^0(x) = \{x^0(0; x), x^0(1; x), \dots, x^0(N; x)\} \quad (1.12)$$

and the value function

$$V_N^0(x) = V_N(x, \mathbf{u}^0(x)) \quad (1.13)$$

(the single argument x denotes the initial state is x at time 0, so $x^0(0; x) = x$). The first control $u^0(0; x)$ in the optimizing sequence $\mathbf{u}^0(x)$ is applied to the plant. The implicit model predictive control law is, therefore,

$$\kappa_N(x) := u^0(0; x) \quad (1.14)$$

and is time invariant. Dynamic Programming could, in principle, be used to determine a sequence $\{V_j(\cdot)\}$ of value functions and a sequence of control laws $\{\kappa_j(\cdot)\}$ where j is time-to-go. Because the optimal control problem is deterministic, the value function $V_N^0(\cdot)$ and its associated control law $\kappa_N(\cdot)$ obtained via dynamic programming are *identical* to the value function and control law in (1.13) and (1.14). It would be preferable to pre-compute $\kappa_N(\cdot)$ using dynamic programming. Since this is usually impossible, MPC computes at event (x, k) the optimal control action $\kappa_N(x)$ rather than pre-computing the control law $\kappa_N(\cdot)$; the difference between using dynamic programming and MPC is, therefore, purely one of implementation. The theoretical properties of MPC derive from the fact that a receding horizon control law $\kappa_N(\cdot)$ is employed rather than, as is conventional, an infinite horizon control law.

In the sequel we find it convenient to refer to the infinite horizon value function $V_\infty^0(\cdot)$ and the associated infinite horizon control law $\kappa_\infty(\cdot)$ for the problem $\mathcal{P}_\infty(x)$ defined as in (1.9)–(1.11) with N replaced by ∞ (under the assumption that the minimum in (1.9) exists).

1.3. Purpose of this review

The discussion immediately above emphasizes the fact that the essential theoretical issues in model predictive control arise from the employment of a receding horizon control law. Stability is an overriding requirement and much of the literature (Section 2) has been devoted to finding sufficient conditions for stability, resulting in a plethora of proposals for model predictive control. Our first, and major, objective (Section 3) is to distill from this large literature the essential ‘ingredients’ required for

stability enabling us to summarize succinctly the major formulations of MPC. In this task, we follow two important streams in the literature, the first a direct approach that employs the value function $V_N^0(\cdot)$ (for a fixed horizon N) as a Lyapunov function, and the second an approach that employs a monotonicity property of a *sequence* $\{V_j^0(\cdot)\}$ of value functions. The theory of stability of model predictive control has reached a relatively mature stage; one purpose of this review is to encourage researchers in nonlinear control to contribute to this area.

Our next concern (Section 4) is robustness, because most other topics, such as output feedback, disturbance attenuation, tracking and adaptation, involve uncertainty. Here progress is limited for several reasons that we discuss. Our main objective is to summarize progress and delineate unresolved difficulties.

We proceed (Section 5) with a brief discussion of various issues, such as tracking, output feedback, and adaptive model predictive control where research has not matured to the same stage as that reached in the study of stability and close our review (Section 6) with some suggestions for future research.

The focus of this review is model predictive control of *constrained* dynamic systems, both linear and nonlinear, since it is the ability of MPC to handle constraints that makes it so attractive to industry; we also give limited attention to the use of model predictive control for problems, such as the control of unconstrained nonlinear systems and time-varying systems, that are otherwise difficult to solve. We do not review, except in passing, the extensive literature on predictive control of unconstrained linear systems, accepting the argument by Bitmead, Gevers and Wertz (1990) that this problem is well handled by linear, quadratic, Gaussian control. We also do not review literature that does not employ state models; this includes significant parts of the process control literature where impulse and step response models are employed and much of the literature on generalized predictive control where transfer function models are employed. Our reason is simple; even if the system being controlled is linear, the presence of constraints makes the controller nonlinear so that the natural tool for establishing stability is Lyapunov theory. Nor do we review applications; an excellent review of these appears in Qin and Badgwell (1997).

MPC literature is somewhat unusual in that, for a while, there were several relatively independent streams: that dealing with theoretical foundations, the ‘process control literature’ that is largely responsible for MPC’s wide-scale adoption by industry, and the literature on generalized predictive control that has its roots in minimum variance and adaptive control. To provide some appreciation, rather than an exhaustive review, of this literature we briefly discuss some aspects of it under these headings in the following section, and follow this with an enumeration of some recent contributions to

model predictive control of constrained systems; these form a basis for our subsequent review.

Because our review is sharply focussed, we mention review papers and books that give a wider background. Earlier reviews include Richalet, Rault, Testud and Papon (1978), García, Prett and Morari (1989), Rawlings, Meadows and Muske (1994), Mayne (1995, 1997), Lee and Cooley (1997), Qin and Badgwell (1997), Chen and Allgöwer (1998a) and Morari and Lee (1999). The review by García et al. (1989) has been particularly influential. Also useful are proceedings of conferences dealing with model predictive control: Clarke (1994), Kantor, García and Carnahan (1997), and Allgöwer and Zheng (1999). The book by Bitmead et al. (1990) gives an excellent, critical exposition of generalized predictive control of linear, unconstrained systems and presents a useful monotonicity argument for establishing stability. Soeterboek (1992) gives a comprehensive exposition of generalized predictive control and its relationship with model predictive control. Mosca (1994) devotes a chapter to predictive control of linear unconstrained systems. Camacho and Bordons (1998), aimed mainly at practitioners, focuses on implementation issues for predictive control.

2. Historical notes

2.1. Theoretical foundations

Obviously relevant to the development of MPC is the literature dealing with the existence of solutions to optimal control problems, characterization of optimal solutions (necessary and sufficient conditions of optimality), Lyapunov stability of the optimally controlled system, and algorithms for the computation of optimal feedback controllers (where possible) and optimal open-loop controls. These topics are dealt with by, for example, Lee and Markus (1967) and Fleming and Rishel (1975).

There are several seminal ideas in the optimal control literature that underly MPC. The first links two major themes in the control revolution of the 1960s: Hamilton–Jacobi–Bellman theory (Dynamic Programming), which provides sufficient conditions for optimality and a constructive procedure for determining an optimal *feedback* controller $u = \kappa(x)$, and the maximum principle, which provides necessary conditions of optimality and motivates computational algorithms for the determination of the optimal open-loop control $u^0(\cdot; x)$ for a given initial state x . The link is

$$\kappa(x) = u^0(0; x)$$

for all x , so that optimal feedback control may be obtained by solving an open-loop control problem (for each state x), an obvious fact appreciated at the very commencement of optimal control; it is implicit, for example,

in the *principle of optimality* (Bellman, 1957) and Lee and Markus (1967, p. 423), specifically state: “One technique for obtaining a feedback controller synthesis from knowledge of open-loop controllers is to measure the current control process state and then compute very rapidly for the open-loop control function. The first portion of this function is then used during a short time interval, after which a new measurement of the process state is made and a new open-loop control function is computed for this new measurement. The procedure is then repeated.” This result follows from the deterministic nature of the optimal control problem (feedback is only necessary when uncertainty is present).

A second observation (Kalman, 1960) is that *optimality does not imply stability*, but that, under certain conditions (stabilizability and detectability), *infinite horizon* optimal controllers are stabilizing, an appropriate Lyapunov function for establishing stability being the value function associated with the infinite horizon optimal control problem.

Solving infinite horizon, open-loop, optimal control problems is not usually practical, especially on-line (apart, of course, from standard H_2 and H_∞ control of linear systems), so a relevant concern is the formulation of *receding horizon* open-loop optimal control problems whose solutions provide stabilizing control. Early examples of results in this direction were Kleinman (1970) and Thomas (1975). Kleinman implicitly considers minimum energy control (of a linear system) to the origin in finite time T , and shows the optimal controller is linear, time invariant ($u = Kx$ where K is easily determined from the controllability Grammian over the interval T). He shows the controller $u = Kx$ is stabilizing using as a Lyapunov function $V(x) = x^T P x$ where P is the inverse of the controllability Grammian over the interval $[0, T]$. Thomas obtains the same result by considering a linear quadratic control problem with control cost, no state cost, and an infinite terminal cost that implicitly adds the *stability constraint* $x(T) = 0$ to the optimal control problem. This is achieved by using $M := P^{-1}$ in place of the Riccati variable P and solving a Riccati-like differential equation for M with terminal condition $M(T) = 0$; this choice implicitly constrains the terminal state to be the origin. Stability is a direct result of incorporating the *stability constraint* $x(T) = 0$ in the optimal control problem. The object of these studies was not to advocate MPC (i.e. on-line solution of the open-loop optimal control problem) but rather to provide alternative methods (alternative to infinite horizon linear-quadratic control) for obtaining linear, time-invariant, stabilizing control. There is no need for on-line optimization since the feedback problem is easily solved. Further extensions to these results were provided in Kwon and Pearson (1977) and Kwon, Bruckstein and Kailath (1983) where a more general linear quadratic problem (with both control and state costs) is considered. The associated Riccati

equation is extensively analysed; in both papers the stability constraint $x(T) = 0$ is implicitly imposed (again by employing the terminal condition $M(T) = 0$ where M is the inverse of the normal Riccati variable P). Interestingly, in the latter paper, the system considered is linear and time varying, and the cost quadratic and time varying so the optimal control is now linear and time varying ($u = K(t)x$). For each t , $K(t)$ may be determined by solving a matrix Riccati differential equation over the interval $[t, t + T]$. The authors show that $K(t)$ may alternatively be obtained by computing a scattering matrix via integration of a differential equation forwards, in real time. From the point of view of further progress, the results of Kleinman, Thomas and Kwon et al. are important, because they show that addition of a terminal constraint ensures stability, but limited, because the stability arguments are appropriate only for unconstrained linear systems.

With this literature, all the ingredients were at hand for defining a stabilizing form of MPC for constrained systems. However, the history of the subject took an unexpected direction that was as exceptional as it was successful.

2.2. The process control literature

When MPC was first advocated by Richalet, Rault, Testud and Papon (1976) for process control, several proposals for MPC had already been made, such as that due to Lee and Markus quoted above, and, even earlier, a proposal, by Propoi (1963), of a form of MPC, using linear programming, for linear systems with hard constraints on control. However, the early proponents of MPC for process control proceeded independently, addressing the needs and concerns of industry. Existing techniques for control design, such as linear quadratic control, were not widely used, perhaps because they were regarded as addressing inadequately the problems raised by constraints, nonlinearities and uncertainty. The applications envisaged were mainly in the petro-chemical and process industries, where economic considerations required operating points (determined by solving linear programmes) situated on the boundary of the set of operating points satisfying all constraints. The dynamic controller therefore has to cope adequately with constraints that would otherwise be transgressed even with small disturbances. The plants were modelled in the early literature by step or impulse responses. These were easily understood by users and facilitated casting the optimal control and identification problems in a form suitable for existing software.

Thus, IDCOM (identification and command), the form of MPC proposed in Richalet et al. (1976,1978), employs a finite horizon pulse response (linear) model, a quadratic cost function, and input and output constraints. The model permits linear estimation, using least squares. The

algorithm for solving the open-loop optimal control problem is a ‘dual’ of the estimation algorithm. As in dynamic matrix control (DMC; Cutler & Ramaker, 1980; Prett & Gillette, 1980), which employs a step response model but is, in other respects, similar, the treatment of control and output constraints is ad hoc. This limitation was overcome in the second-generation program, quadratic dynamic matrix control (QDMC; García & Morshedi, 1986) where quadratic programming is employed to solve exactly the constrained open-loop optimal control problem that results when the system is linear, the cost quadratic, and the control and state constraints are defined by linear inequalities. QDMC also permits, if required, temporary violation of some output constraints, effectively enlarging the set of states that can be satisfactorily controlled. The third generation of MPC technology, introduced about a decade ago, “distinguishes between several levels of constraints (hard, soft, ranked), provides some mechanism to recover from an infeasible solution, addresses the issues resulting from a control structure that changes in real time, and allows for a wider range of process dynamics and controller specifications” (Qin & Badgwell, 1997). In particular, the Shell multivariable optimizing control (SMOC) algorithm allows for state-space models, general disturbance models and state estimation via Kalman filtering (Marquis & Broustail, 1988). The history of the three generations of MPC technology, and the subsequent evolution of commercial MPC, is well described in the last reference. The substantial impact that this technology has had on industry is confirmed by the number of applications (probably exceeding 2000) that make it a multi-million dollar industry.

The industrial proponents of MPC did not address stability theoretically, but were obviously aware of its importance; their versions of MPC are not automatically stabilizing. However, by restricting attention to stable plants, and choosing a horizon large compared with the ‘settling’ time of the plant, stability properties associated with an infinite horizon are achieved. Academic research, stimulated by the unparalleled success of MPC, commenced a theoretical investigation of stability. Because Lyapunov techniques were not employed initially, stability had to be addressed within the restrictive framework of linear analysis, confining attention to model predictive control of linear unconstrained systems. The original finite horizon formulation of the optimal control problem (without any modification to ensure stability) was employed. Researchers therefore studied the effect of control and cost horizons and cost parameters on stability when the system is linear, the cost quadratic, and hard constraints are absent. See García et al. (1989) for an excellent review of this literature. A typical result establishes the existence of finite control and cost horizons such that the resultant model predictive controller is stabilizing.

2.3. Generalized predictive control

Largely independently a closely related methodology developed in the adaptive control literature. In this literature models of the form

$$A(q^{-1})y(t) = B(q^{-1})u(t) + C(q^{-1})e(t),$$

where y is the output, u the control, and e a white-noise disturbance and A , B and C are polynomials in the backward shift operator q^{-1} , were employed because of their greater generality than impulse or step response models. Deficiencies in minimum variance control (Åström, 1970) (a precursor of generalized predictive control with horizon $N = 1$) motivated an extension of the horizon yielding a formulation (De Keyser & Van Cauwenberghe, 1979; Ydstie, 1984; Peterka, 1984) and (Clarke, Mohtadi & Tuffs, 1987a,b) that, in its deterministic version, is very close to that of MPC. Because of the finite horizon, stability was not guaranteed in the original versions of generalized predictive control and was achieved, as in early versions of MPC, by tuning cost and horizon parameters, an approach criticized in Bitmead et al. (1990) as ‘playing games’ instead of enforcing stability by modification of the optimal control problem. This criticism was supported by examples presented in Soeterboek (1992). At the commencement of the 1990s, stability of predictive controllers was successfully addressed in a series of papers (Mosca, Lemos & Zhang, 1990; Clarke & Scattolini, 1991; Mosca & Zhang, 1992); these papers established stability of linear, *unconstrained*, input/output systems (systems for which the state is not accessible). Stability is achieved by imposing terminal equality constraints on inputs and outputs over a finite interval; this constraint is equivalent to the terminal constraint employed by Kleinman and Kwon et al. if the interval is long enough and observability conditions are satisfied. Because the system is linear, the cost quadratic, the terminal constraint a linear equality, and control and state constraints are absent, the optimal control may be computed algebraically. The early literature on generalized predictive control is an important strand in the development of the subject but, because it does not deal with control and state constraints, falls outside the scope of this review.

2.4. The evolution of model predictive control

Early versions of MPC and generalized predictive control did not automatically ensure stability, thus requiring tuning. It is therefore not surprising that research in the 1990s devoted considerable attention to this topic. Indeed, concern for stability has been a major engine for generating different formulations of MPC. In time, differences between model predictive, generalized predictive, and receding horizon control became irrelevant; we therefore use MPC as a generic title in the sequel for that

mode of control in which the current control action is determined by solving on-line an optimal control problem. Our purpose here (Section 2.4) is not to review comprehensively recent literature but merely to introduce a few significant advances made in the 1990s that motivate the analysis presented in Section 3.

2.4.1. Stability analysis

Model predictive control of constrained systems is nonlinear necessitating the use of Lyapunov stability theory, a tool neglected in the MPC literature during the 1970s and the early 1980s with a notable exception: Chen and Shaw (1982) showed, in a paper unfortunately not noticed for some time, that the value function (of a finite horizon optimal control problem) could be used as Lyapunov function to establish stability of continuous time receding horizon control of unconstrained systems when a terminal equality constraint is employed. These results on continuous-time systems were not extended until Mayne and Michalska (1990). Keerthi and Gilbert (1988) first employed the value function as a Lyapunov function for establishing stability of model predictive control of time-varying, *constrained, nonlinear*, discrete-time systems (when a terminal equality constraint is employed); thereafter, the value function was almost universally employed as a natural Lyapunov function for stability analysis of model predictive control.

2.4.2. Stabilizing modifications to $\mathcal{P}_N(x)$

The 1990s saw many proposals for modifying the open-loop optimal control problem $\mathcal{P}_N(x)$ employed in model predictive control of constrained and/or nonlinear systems so that closed-loop stability could be guaranteed. We list these proposals below, deferring a more detailed discussion until Section 3. The modifications correspond, in the main, to proposals for the terminal cost $F(\cdot)$ and the terminal constraint set X_f . Reference should be made to Section 1.2 where the open-loop optimal control problem $\mathcal{P}_N(x)$ is defined. To simplify presentation, we give discrete-time analogues of results derived for continuous-time systems.

2.4.2.1. Terminal equality constraint. In this version of model predictive control, the terminal cost $F(\cdot)$ and the terminal constraint $x(N) \in X_f$ in the open-loop optimal control problem $\mathcal{P}_N(x)$ satisfy $F(x) \equiv 0$ and $X_f = \{0\}$. The first proposal for this form of model predictive control for time-varying, *constrained, nonlinear*, discrete-time systems was made by Keerthi and Gilbert (1988). This paper is particularly important; it provides a definitive stability analysis of this version of discrete-time receding horizon control (under mild conditions of controllability and observability) and shows the value function $V_N^0(x)$ associated with the finite horizon optimal control problem approaches that of the infinite horizon problem as the horizon approaches infinity. This paper remains

a key reference on the stabilizing properties of model predictive control and subsumes much of the later literature on discrete-time MPC that uses a terminal equality constraint. It is an interesting fact that model predictive control with a terminal equality constraint can also be employed to stabilize systems that cannot be stabilized by continuous feedback controllers (Meadows, Henson, Eaton & Rawlings, 1995).

Continuous-time versions are described in Chen and Shaw (1982) and Mayne and Michalska (1990). By defining a non-minimal state in terms of current and past inputs and outputs, Chisci and Mosca (1994) establish stability of model predictive control of linear unconstrained systems of the form $A(q^{-1})y(t) = B(q^{-1})u(t)$ where A and B are polynomials in the delay operator q^{-1} ; the terminal constraint $x(N) = 0$ translates to constraints on $u(t)$ and $y(t)$ over the interval $N, N + 1, \dots, N + n - 1$, where n is the state dimension. A similar constraint is employed in Bemporad, Chisci and Mosca (1995) for linear unconstrained systems of the form $A(q^{-1})y(t) = B(q^{-1})u(t) + C(q^{-1})e(t)$ where $\{e(t)\}$ is a white-noise process (the cost is defined as a conditional expectation). More relevant to this review are the papers by De Nicolao, Magni and Scattolini (1996a) and Magni and Sepulchre (1997) that employ the terminal constraint $x(N) = 0$ to establish closed-loop stability when the system is nonlinear and unconstrained.

2.4.2.2. Terminal cost function. One of the earliest proposals for modifying $\mathcal{P}_N(x)$ to ensure closed-loop stability was the addition of a terminal cost. In this version of model predictive control, the terminal cost $F(\cdot)$ is non-trivial and there is no terminal constraint so that $X_f = \mathbb{R}^n$. The proposal was made (Bitmead et al., 1990) in the context of predictive control of unconstrained linear system for which the choice $F(x) = (1/2)x^T P_f x$ is appropriate. The matrix P_f , the terminal value of the Riccati difference equation, is chosen so that the sequence $\{P_j\}$, obtained by solving the Riccati difference equation in reverse time (j is time-to-go) with terminal condition $P_0 = P_f$ is monotonically non-increasing (i.e. $P_{j+1} \leq P_j$ for all $j \geq 0$). Using the fake algebraic Riccati equation (FARE), the authors show (under mild conditions) that this choice yields a value function $V_N^0(x) = (1/2)x^T P_N x$ and a receding horizon controller $\kappa_N(x) = K_N x$ that is stabilizing. The methodology for establishing stability is interesting and can be extended to nonlinear systems; we discuss this more fully in Section 3. The controller $\kappa_N(x) = K_N x$ obtained by Bitmead et al. may (and should) be computed off-line.

A more interesting example is a proposal (Rawlings & Muske, 1993) that stimulated much subsequent research. When the system being controlled is linear and stable, and subject to a convex control constraint, the terminal cost $F(\cdot)$ is chosen to be the (infinite horizon) value function associated with the (stabilizing) controller

$u = \kappa_f(x) \equiv 0$. Gauthier and Bornard (1983) similarly use a local controller $\kappa_f(x) \equiv 0$ to convert an infinite horizon, open-loop optimal control problem to a finite horizon problem and establish stability when the system is linear and unconstrained. These results are equivalent to having an infinite cost horizon and a finite control horizon. In interesting related research, Alamir and Bornard (1995) employ a control horizon N and a cost horizon M and show that closed-loop stability ensues if M is large enough.

Can this technique for achieving stability (by adding only a terminal cost) be successfully employed for constrained and/or nonlinear systems? From the literature the answer may appear affirmative. However, in this literature there is an implicit requirement that $x(N) \in X_f$ is satisfied for every initial state x in a given compact set, and this is automatically satisfied if N is chosen sufficiently large; the constraint $x(N) \in X_f$ then need not be included explicitly in the optimal control problem actually solved on-line. Whether this type of model predictive control is regarded as having only a terminal cost or having both a terminal cost and a terminal constraint is a matter of definition; we prefer to consider it as belonging to the latter category as the constraint is necessary even though it is automatically satisfied if N is chosen sufficiently large.

2.4.2.3. Terminal constraint set. In this version of model predictive control, X_f is a subset of \mathbb{R}^n containing a neighbourhood of the origin and $F(x) \equiv 0$ (no terminal cost). The purpose of the model predictive controller is to steer the state to X_f in finite time. Inside X_f , a local stabilizing controller $\kappa_f(\cdot)$ is employed; this form of model predictive control is therefore sometimes referred to as dual mode, and was proposed, in the context of constrained, continuous-time, nonlinear systems, in Michalska and Mayne (1993) where a variable horizon N is employed. Fixed horizon versions for constrained, nonlinear, discrete-time systems are proposed in Chisci, Lombardi and Mosca (1996) and Sokaert, Mayne and Rawlings (1999). Dual mode control is also employed, in a different context, in Szaiaer and Damborg (1990).

2.4.2.4. Terminal cost and constraint set. Most recent model predictive controllers belong to this category. There are a variety of good reasons for incorporating both a terminal cost and a terminal constraint set in the optimal control problem $\mathcal{P}_N(x)$. Ideally, the terminal cost $F(\cdot)$ should be the infinite horizon value function $V_\infty^0(\cdot)$; if this were the case, then $V_N^0(\cdot) = V_\infty^0(\cdot)$, on-line optimization would be unnecessary, and the known advantages of an infinite horizon, such as stability and robustness, would automatically accrue. Nonlinearity and/or constraints render this impossible, but it is possible to choose $F(\cdot)$ so that it is exactly or approximately equal to $V_\infty^0(\cdot)$ in a suitable neighbourhood of the origin. Choosing

X_f to be an appropriate subset of this neighbourhood yields many advantages and motivates the choice of $F(\cdot)$ and X_f in most of the examples of this form of model predictive control.

Perhaps the earliest proposal for this form of model predictive control is the remarkably prescient paper (Sznaier & Damborg, 1987). The system considered is linear ($f(x, u) = Ax + Bu$) and the state and control constraint sets, \mathbb{X} and \mathbb{U} respectively, are polytopes. Then $F(x)$ and $\kappa_f(x) = K_f x$ are chosen to be the value function and optimal controller for the *unconstrained* infinite horizon problem; this is a standard *LQ* problem so $F(\cdot)$ and $\kappa_f(\cdot)$ are easily computed. The terminal constraint set X_f is chosen to be the *maximal output admissible set* (Gilbert & Tan, 1991) of the system $x^+ = (A + BK_f)x$. When the current state is x , the horizon $N = N_x$ can be chosen so that the terminal state $x^0(N; x, 0)$ (corresponding to the optimal control $\mathbf{u}^0(x)$) lies in X_f ; hence $x^0(N; x) \in X_f$ even though $\mathcal{P}_N(x)$ does *not* include the terminal constraint $x(N) \in X_f$. With this choice of N , $V_N^0(x) = V_\infty^0(x)$ and $\kappa_N(x) = \kappa_\infty(x)$; stability results as proven in Chmielewski and Manousiouthakis (1996), and Scokaert and Rawlings (1998). In this literature, the set X_f is introduced as a consequence of the control and state constraints that render determination of $V_\infty^0(x)$ impossible (except in X_f).

For the case when the system is nonlinear but there are no state or control constraints, Parisini and Zoppoli (1995) and Jadbabaie, Yu and Hauser (1999) (in the context of continuous-time systems) use a stabilizing local control law $\kappa_f(\cdot)$, a terminal cost function $F(\cdot)$ that is a (local) Lyapunov function for the stabilized system, and a terminal constraint set X_f that is a level set of $F(\cdot)$ and is positively invariant for the system $x^+ = f(x, \kappa_f(x))$. The terminal constraint is *omitted* from the problem $\mathcal{P}_N(x)$ solved on-line but it is nevertheless shown that this constraint is automatically satisfied for all initial states in a level set of $V_N^0(\cdot)$; the resultant closed-loop system is asymptotically (or exponentially) stabilizing with a region of attraction that is this level set of $V_N^0(\cdot)$.

When the system is both nonlinear and constrained, $F(\cdot)$ and X_f include features from both of the two examples immediately above. In both De Nicolao, Magni and Scattolini (1996c) and Chen and Allgöwer (1998b) (modulo the fact that continuous-time systems are considered in the latter paper) $\kappa_f(\cdot)$ is chosen to stabilize the linearized system $x^+ = Ax + Bu$, where $A := f_x(0, 0)$ and $B := f_u(0, 0)$. Allgöwer employs a quadratic and De Nicolao et al. a non-quadratic terminal cost $F(\cdot)$. Both employ a terminal constraint set X_f that is positively invariant for the nonlinear system $x^+ = f(x, \kappa_f(x))$ and that satisfies $X_f \subset \mathbb{X}$ and $\kappa_f(X_f) \subset \mathbb{U}$. Chen and Allgöwer refer to their approach as ‘quasi-infinite horizon predictive control’ because the finite horizon optimal control problem they employ approximates the full infinite horizon problem.

2.5. Conclusion

The discussion above reveals the presence of several ‘ingredients’ that have been found useful in developing stabilizing model predictive controllers; these ingredients are a terminal cost $F(\cdot)$, a terminal constraint set X_f , and a local controller $\kappa_f(\cdot)$. In Section 3, we present conditions on $F(\cdot)$, X_f , and $\kappa_f(\cdot)$ that, if satisfied, ensure that the model predictive controller is stabilizing and illustrate these conditions by showing that they are satisfied in most of the controllers presented above.

3. Stability and inverse optimality

3.1. Notation

For any function $\phi(x)$, let $\overset{*}{\phi}(x, u)$ denote the change in $\phi(\cdot)$ as the state changes from x to $x^+ = f(x, u)$, i.e.

$$\overset{*}{\phi}(x, u) := \phi(f(x, u)) - \phi(x).$$

$\overset{*}{\phi}(x, u)$ is the analogue of $\dot{\phi}(x, u)$ in the continuous-time case (when the system satisfies $\dot{x} = f(x, u)$).

3.2. Introduction

Two distinct, but related, methods for establishing stability have evolved in the literature, each yielding its own insights. Both approaches employ the value function $V_N^0(\cdot)$ as a Lyapunov function. We assume, in the sequel, that the value function is continuous. The first approach, which we will call the direct method, employs the value function and obtains conditions on $F(\cdot)$, X_f , and $\kappa_f(\cdot)$ that ensure

$$\overset{*}{V}_N^0(x, \kappa_N(x)) + \ell(x, \kappa_N(x)) \leq 0. \quad (3.1)$$

In this methodology, employed *inter alia* in Keerthi and Gilbert (1988), Mayne and Michalska (1990), Rawlings and Muske (1993), Michalska and Mayne (1993), Chmielewski and Manousiouthakis (1996), Scokaert and Rawlings (1998) and Chen and Allgöwer (1998), computation of $V_N^0(x^+)$, $x^+ := f(x, \kappa_N(x))$ is avoided by computing an upper bound $V_N(x^+, \tilde{\mathbf{u}}(x))$ for $V_N^0(x^+)$ using a feasible control $\tilde{\mathbf{u}}(x)$ for $\mathcal{P}_N(x^+)$.

The second method uses the fact that

$$\overset{*}{V}_N^0(x, \kappa_N(x)) + \ell(x, \kappa_N(x)) = V_N^0(x^+) - V_{N-1}^0(x^+)$$

and shows the right-hand side of this equation is negative, either directly or by showing $V_1^0(\cdot) \leq V_0^0(\cdot)$ and exploiting monotonicity ($V_1^0(\cdot) \leq V_0^0(\cdot)$ implies $V_{i+1}^0(\cdot) \leq V_i^0(\cdot)$ for all $i \geq 0$). This approach has been used *inter alia* in Chen and Shaw (1982), Bitmead et al.

(1990), Bemporad et al. (1995), De Nicolao et al. (1996a,1996c), Magni and Sepulchre (1997), De Nicolao and Sepulchre (1997), Meadows (1997) and Primbs and Nevistić (1997). These two approaches are reviewed in Allgöwer and Zheng (1999).

3.3. Direct method

Our purpose here is to distill from the extensive literature on model predictive control essential principles, presented below as axioms, that ensure closed-loop stability. This requires the determination of appropriate conditions on the ingredients $F(\cdot)$, X_f and $\kappa_f(\cdot)$ present in most forms of model predictive control. For each integer k , let X_k denote the set of states x steerable, by admissible control sequences to X_f in k steps or less. An admissible (or feasible) control sequence $\mathbf{u} = \{u(0), u(1), \dots, u(k-1)\}$ satisfies the control, state and terminal constraints, i.e. $u(i) \in \mathbb{U}$ for $i = 0, 1, \dots, k-1$, $x^u(i; x) \in \mathbb{X}$ for $i = 0, 1, \dots, k$ and $x^u(k; x) \in X_f$. The set of states that can be controlled by model predictive control with fixed horizon N is X_N . Suppose then, that $x \in X_N$, and that the control sequence $\mathbf{u}^0(x)$ that solves $\mathcal{P}_N(x)$ has been determined. Let $\mathbf{x}^0(x) = \{x, x^0(1; x), \dots, x^0(N; x)\}$ denote the optimal state trajectory. The model predictive control $u = \kappa_N(x) = u^0(0; x)$ steers the initial state x to the successor state $x^+ = x^0(1; x) = f(x, \kappa_N(x))$. We wish to determine a feasible control sequence $\tilde{\mathbf{u}}(x)$ for x^+ and hence, an upper bound for $V_N^0(x^+)$. Since $\{u^0(0; x), u^0(1; x), \dots, u^0(N-1; x)\}$ is a feasible control sequence for $\mathcal{P}_N(x)$, $x^0(N; x) \in X_f$, and the (abbreviated) control sequence $\{u^0(1; x), \dots, u^0(N-1; x)\}$ steers x^+ to $x^0(N; x) \in X_f$. To obtain a feasible control for $\mathcal{P}_N(x^+)$, we add one further element u to this sequence, obtaining $\{u^0(1; x), \dots, u^0(N-1; x), u\}$; this sequence will be feasible for $\mathcal{P}_N(x^+)$ if $u \in \mathbb{U}$ and u steers $x^0(N; x) \in X_f$ to a state $f(x^0(N; x), u) \in X_f$ which is the case if $u = \kappa_f(x^0(N; x))$ and X_f and the local controller $\kappa_f(\cdot)$ have the following properties:

$$X_f \subset \mathbb{X}, \kappa_f(x) \in \mathbb{U} \quad \text{and} \quad f(x, \kappa_f(x)) \in X_f \quad \forall x \in X_f$$

so that X_f is positively invariant when the control law is $\kappa_f(\cdot)$. If these conditions are satisfied, the control sequence

$$\tilde{\mathbf{u}}(x) := \{u^0(1; x), \dots, u^0(N-1; x), \kappa_f(x^0(N; x))\}$$

is feasible for $\mathcal{P}_N(x^+)$. The state trajectory resulting from initial state $x^+ = x^0(1; x)$ and control sequence $\tilde{\mathbf{u}}(x)$ is

$$\tilde{\mathbf{x}}(x) := \{x^0(1; x), \dots, x^0(N; x), f(x^0(N; x), \kappa_f(x^0(N; x)))\}.$$

The associated cost is

$$\begin{aligned} V_N(x^+, \tilde{\mathbf{u}}(x)) &= V_N^0(x) - \ell(x, \kappa_N(x)) - F(x^0(N; x)) \\ &\quad + \ell(x^0(N; x), \kappa_f(x^0(N; x))) \\ &\quad + F(f(x^0(N; x), \kappa_f(x^0(N; x)))) \end{aligned}$$

This cost, which is an upper bound for $V_N^0(x^+)$, satisfies

$$V_N(x^+, \tilde{\mathbf{u}}(x)) \leq V_N^0(x) - \ell(x, \kappa_N(x))$$

$$\text{if } F(x, \kappa_f(x)) + \ell(x, \kappa_f(x)) \leq 0, \text{ i.e. if}$$

$$[F + \ell](x, \kappa_f(x)) \leq 0 \quad \forall x \in X_f$$

since then the sum of the last three terms in the expression for $V_N(x^+, \tilde{\mathbf{u}}(x))$ is less than or equal to zero; this condition holds if $F(\cdot)$ is a control Lyapunov function in a neighbourhood of the origin (i.e. the infimum of $F(f(x, u) - F(x))$ with respect to u is less than zero in this neighbourhood) and $\kappa_f(\cdot)$ and X_f are chosen appropriately. If this condition holds, then (3.1) also holds for all $x \in X_N$ which is sufficient to ensure that the state of the closed-loop system $x^+ = f(x, \kappa_N(x))$ converges to zero as $k \rightarrow \infty$ if its initial state lies in X_N . This motivates the following conditions that, if satisfied, ensure closed-loop asymptotic (exponential) stability if further minor assumptions are satisfied:

- A1: $X_f \subset \mathbb{X}$, X_f closed, $0 \in X_f$ (state constraint satisfied in X_f).
- A2: $\kappa_f(x) \in \mathbb{U}$, $\forall x \in X_f$ (control constraint satisfied in X_f).
- A3: $f(x, \kappa_f(x)) \in X_f$, $\forall x \in X_f$ (X_f is positively invariant under $\kappa_f(\cdot)$).
- A4: $[F + \ell](x, \kappa_f(x)) \leq 0$, $\forall x \in X_f$ ($F(\cdot)$ is a local Lyapunov function).

Assumption A4 implies A3 if X_f is a level set of $F(\cdot)$; this is a common situation. These conditions provide a concise characterization, as we show in Section 3.7, of the model predictive controllers described earlier. The control sequence $\tilde{\mathbf{u}}(x)$, in addition to providing an upper bound for $V_N^0(x)$, can also be usefully employed to initialize the algorithm for solving $\mathcal{P}_N(x^+)$. The conditions are, of course, merely sufficient.

3.4. Monotonicity of $\{V_j^0(\cdot)\}$

To introduce this approach, we note, as observed by Chen and Shaw (1982), that, by the principle of optimality with $x^+ := f(x, \kappa_N(x))$:

$$V_N^0(x) = \ell(x, \kappa_N(x)) + V_{N-1}^0(x^+)$$

for all $x \in X_N$ so that

$$V_N^0(x) = \ell(x, \kappa_N(x)) + V_N^0(x^+) + [V_{N-1}^0(x^+) - V_N^0(x^+)]$$

and

$$V_N^0(x, \kappa_N(x)) + \ell(x, \kappa_N(x)) = [V_N^0(x^+) - V_{N-1}^0(x^+)]. \quad (3.2)$$

Hence the desired inequality (3.1) is recovered if $V_N^0(x) \leq V_{N-1}^0(x)$ for all $x \in X_{N-1}$. This, if axioms A1–A4 hold, can be shown to be true in much the same way as was employed in Section 3.2; an upper bound for $V_N^0(x^+)$ can be employed by taking the (optimizing) control sequence (consisting of $N - 1$ controls) for $\mathcal{P}_{N-1}^0(x^+)$ and adding the control $\kappa_f(x^0(N-1; x^+)) = \kappa_f(x^0(N; x))$; the resultant sequence is precisely $\tilde{u}(x)$ defined above and is feasible for $\mathcal{P}_{N-1}^0(x^+)$. Hence

$$\begin{aligned} V_N(x^+, \tilde{u}(x)) &= V_{N-1}^0(x^+) + \ell(x^0(N; x), \kappa_f(x^0(N; x))) \\ &\quad - F(x^0(N; x)) + F(f(x^0(N; x), \kappa_f(x^0(N; x)))) \\ &\geq V_N^0(x^+) \end{aligned}$$

which, by A4, implies $V_N^0(x^+) \leq V_{N-1}^0(x^+)$; this inequality and (3.2) establish the desired inequality (3.1). This procedure for establishing stability was used, for example, by De Nicolao et al. (1996a) for unconstrained nonlinear systems, by Magni and Sepulchre (1997) for unconstrained, nonlinear, continuous-time systems and by De Nicolao et al. (1996c) for constrained, nonlinear systems. The same result can be obtained using monotonicity. Bitmead et al. (1990) show, in the context of linear quadratic (unconstrained) problems where $V_i(x) = x^T P_i x$, that, if $V_1^0(\cdot) \leq V_0^0(\cdot)$, then $V_{i+1}^0(\cdot) \leq V_i^0(\cdot)$ for all $i \geq 0$; this is the *monotonicity* property that, for *constrained* nonlinear systems, may be expressed as follows:

Suppose $V_1^0(x) \leq V_0^0(x)$ for all $x \in X_0 = X_f$. Then $V_{i+1}^0(x) \leq V_i^0(x)$ for all $x \in X_i$, all $i \geq 0$.

To prove this assertion, assume $V_i^0(x) \leq V_{i-1}^0(x)$ for all $x \in X_{i-1}$. Then

$$\begin{aligned} V_{i+1}^0(x) - V_i^0(x) &= \ell(x, \kappa_i(x)) + V_i^0(f(x, \kappa_i(x))) \\ &\quad - \ell(x, \kappa_i(x)) - V_{i-1}^0(f(x, \kappa_i(x))) \\ &\leq 0 \quad \text{for all } x \in X_i, \end{aligned}$$

where we make use of the fact that $f(x, \kappa_i(x)) \in X_{i-1}$ if $x \in X_i$. Induction establishes the assertion. In view of the monotonicity property it is merely necessary to establish that $V_1^0(x) \leq V_0^0(x)$ for all $x \in X_0 = X_f$:

Suppose $F(\cdot)$, X_f and $\kappa_f(\cdot)$ satisfy A1–A4. Then $V_1^0(x) \leq V_0^0(x)$ for all $x \in X_0 = X_f$.

This may be proven as follows:

$$\begin{aligned} V_1^0(x) &= \ell(x, \kappa_1(x)) + V_0^0(f(x, \kappa_1(x))) \\ &\quad \text{(by principle of optimality)} \end{aligned}$$

$$\begin{aligned} &\leq \ell(x, \kappa_f(x)) + V_0^0(f(x, \kappa_f(x))) \\ &\quad \text{(by optimality of } \kappa_1(\cdot) \text{)} \\ &= \ell(x, \kappa_f(x)) + F(f(x, \kappa_f(x))) \\ &\quad \text{(by definition of } V_0^0(\cdot) \text{)} \\ &\leq F(x) = V_0^0(x) \quad \text{(by A4).} \end{aligned}$$

It then follows that $V_N^0(x) \leq V_{N-1}^0(x)$ for all $x \in X_{N-1}$ if A1–A4 are satisfied; this establishes the desired inequality (3.1) and, hence, asymptotic (exponential) stability of the closed-loop system.

3.5. Inverse optimality

Bitmead et al. (1990) show, in the context of linear, unconstrained systems, that if the monotonicity property holds, the (finite horizon) value function $V_N^0(\cdot)$ is also the *infinite horizon* value function of a modified problem, an interesting example of inverse optimality. The known advantages of infinite horizon optimal control then accrue. Of course, stability can be established independently, as shown above. Inverse optimality can be easily extended to the nonlinear case as shown in Magni and Sepulchre (1997) for nonlinear, unconstrained continuous-time systems. Here we deal with the nonlinear, discrete-time case. Eq. (3.2) may be rewritten in the form

$$V_N^0(x) = \bar{\ell}(x, \kappa_N(x)) + V_N^0(f(x, \kappa_N(x))) \tag{3.3}$$

for all $x \in X_N$ where

$$\bar{\ell}(x, u) := \ell(x, u) + [V_{N-1}^0 - V_N^0](f(x, \kappa_N(x)))$$

and $\bar{\ell}(x, u) \geq \ell(x, u) \geq c|(x, u)|^2$ if A1–A4 hold. Eq. (3.3) is a fake Hamilton–Jacobi–Bellman algebraic equation, corresponding to the infinite horizon optimal control problem $\mathcal{P}_\infty(x)$ with $\ell(\cdot)$ replaced by $\bar{\ell}(\cdot)$; it is the analogue of the fake Riccati equation introduced by Poubelle, Bitmead and Gevers (1988). The controller $\kappa_N(\cdot)$ is optimal for the modified problem. Hence, model predictive control of unconstrained systems, linear or nonlinear, derived using a finite horizon optimal control problem, inherits the known robustness properties of infinite horizon optimal control as shown in Magni and Sepulchre (1997) for nonlinear systems satisfying $\dot{x} = f(x) + g(x)u$.

3.6. Continuous-time systems

Suppose the system is described by

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

and that

$$V(x, t, u(\cdot)) := \int_t^T \ell(x(s), u(s)) ds + F(x(T)),$$

where $x(s) = x^u(s; x, t)$. The constraints remain the same:

$$x(t) \in \mathbb{X}, \quad u(t) \in \mathbb{U}, \quad x(T) \in X_f.$$

For all $\tau \in [0, T]$, let X_τ denote the set of states that can be steered to X_f , by an admissible control, in time τ . Let $V^0(x, t)$, denote the value function and $\kappa^0(x, t)$ the optimal control law for problem $\mathcal{P}(x, t)$ defined by

$$V^0(x, t) = \min\{V(x, t, \mathbf{u}) \mid \mathbf{u} \in \mathcal{U}_{[t, T]}\},$$

where $\mathcal{U}_{[t, T]}$ is the set of admissible controls defined on $[t, T]$ (satisfying control, state and terminal constraints). Using time invariance, the optimal control problem solved on-line is $\mathcal{P}_T(x) := \mathcal{P}(x, 0)$ and its value function is $V_T^0(\cdot) := V^0(x, 0)$. The optimizing control and state trajectories, assumed to exist, are $u^0(\cdot; x)$ and $x^0(\cdot; x)$ (the argument x indicates the initial state is x at $t = 0$) so that

$$V_T^0(x) := V^0(x, 0) = V(x, 0, u^0(\cdot)).$$

Here the suffix T is the time-to-go whereas t in $V^0(x, t)$ is current time. The implicit model predictive controller is $\kappa_T(x) := u^0(0; x)$ (practically, a variant of model predictive control in which $u^0(\cdot; x)$ is applied for time δ before \mathcal{P}_T is resolved would be employed). Given a function $\phi(\cdot): \mathbb{R}^n \rightarrow \mathbb{R}$, let $\dot{\phi}(x, u)$ denote its directional derivative in the direction $f(x, u)$; if $\phi(\cdot)$ is differentiable

$$\dot{\phi}(x, u) = \phi_x(x)f(x, u).$$

The ingredients $F(\cdot)$, X_f and $\kappa_f(\cdot)$ for the continuous-time case are required to satisfy

- B1: $X_f \subset \mathbb{X}$, X_f closed, $0 \in X_f$.
- B2: $\kappa_f(x) \in \mathbb{U}$, $\forall x \in X_f$.
- B3: X_f is positively invariant for $\dot{x} = f(x, \kappa_f(x))$.
- B4: $[\dot{F} + \ell](x, \kappa_f(x)) \leq 0$, $\forall x \in X_f$.

Condition B4 implies B3 if X_f is a level set of $F(\cdot)$. Conditions B1–B4 ensure

$$[\dot{V}_T + \ell](x, \kappa_T(x)) \leq 0$$

for all $x \in X_T$ and, hence, are sufficient, modulo further modest conditions, for asymptotic (exponential) stability of the closed-loop system.

A monotonicity property, analogous to that for discrete-time systems, may be established under the (strong) assumption that $V^0(\cdot)$ is continuously differentiable (constraints may invalidate this assumption):

Suppose $(\partial/\partial t)V^0(x, T) \geq 0$ for all $x \in X_0 = X_f$. Then $(\partial/\partial t)V^0(x, t) \geq 0$ for all $x \in X_{T-t}$, all $t \in [0, T]$.

This result is a direct consequence of the fact that

$$(\partial/\partial t)V^0(x, t) = (\partial/\partial t)V^0(x^0(s; x, t), s)$$

for all $s, t \in [0, T]$. If $V^0(\cdot)$ is continuously differentiable, the fake H–J equation is (Magni & Sepulchre, 1997)

$$\bar{\ell}(x, \kappa_T(x)) + (\partial/\partial x)V^0(x, 0)f(x, \kappa_T(x)) = 0,$$

where $V^0(x, 0) = V_T(x)$ and

$$\bar{\ell}(x, u) := \ell(x, u) + (\partial/\partial t)V^0(x, 0)$$

so that $\bar{\ell}(x, u) \geq \ell(x, u)$ if $(\partial/\partial t)V^0(x, 0) \geq 0$ for all $x \in X_T$ which is the case (by the monotonicity property) if $(\partial/\partial t)V(x, T) \geq 0$ for all $x \in X_0 = X_f$:

Suppose B1–B4 are true. Then $(\partial/\partial t)V^0(x, T) \geq 0$ for all $x \in X_0 = X_f$.

This may be proven as before. Let $\kappa_0(\cdot) := \kappa^0(\cdot, T)$ denote the optimal control law when the time-to-go is 0, i.e. current time is t . Then:

$$\begin{aligned} -(\partial/\partial t)V^0(x, T) &= \ell(x, \kappa_0(x)) + (\partial/\partial x)V^0(x, T)f(x, \kappa_0(x)) \\ &\leq \ell(x, \kappa_f(x)) + (\partial/\partial x)V^0(x, T)f(x, \kappa_f(x)) \\ &= \ell(x, \kappa_f(x)) + \dot{F}(x, \kappa_f(x)) \leq 0, \end{aligned}$$

where the first line follows from the principle of optimality, the second from the optimality of $\kappa_0(\cdot)$, and the third by B4. If B1–B4 are satisfied,

$$[\dot{V}_T^0 + \bar{\ell}](x, \kappa_T(x)) = 0,$$

where $\bar{\ell}(\cdot) \geq \ell(\cdot)$ from which asymptotic (exponential) stability may be deduced.

3.7. Characterization of model predictive controllers

We showed above that, regardless of whether the direct or monotonicity approach is used, closed-loop asymptotic stability, with a domain of attraction X_N , results if A1–A4 hold. These axioms may be satisfied by a variety of choices, giving rise to many variants of model predictive control. Our purpose here is to show that all the model predictive controllers discussed above satisfy these axioms.

3.7.1. Terminal state

This variant of model predictive control (Keerthi & Gilbert, 1988) is characterized by its terminal constraint $x(N) = 0$ so that $X_f = \{0\}$. A terminal cost is pointless so $F(\cdot)$ is chosen to be zero. The local controller must maintain the state in X_f . The ingredients for this form of model predictive control are, therefore $F(x) \equiv 0$ (no terminal cost), $X_f = \{0\}$ (terminal equality constraint), and $\kappa_f(x) \equiv 0$ (zero control maintains the state at the origin). The functions $F(\cdot)$ and $\kappa_f(\cdot)$ need only be defined on X_f , i.e. at $x = 0$. Satisfaction of the axioms is easily checked. Since $X_f = \{0\} \in \mathbb{X}$, A1 is satisfied. Also

$\kappa(0) = 0 \in \mathbb{U}$ so that A2 is satisfied. A3 follows from the fact that $f(0, \kappa_f(0)) = f(0, 0) = 0 \in X_f$. Finally,

$$^* F(0, \kappa_f(0)) + \ell(0, \kappa_f(0)) = 0$$

since $\kappa_f(0) = 0, f(0, 0) = 0$ and $\ell(0, 0) = 0$, ensuring satisfaction of A4. Asymptotic (exponential) stability with a region of attraction X_N results.

3.7.2. Terminal cost

This variant of model predictive control is characterized by the use of a terminal cost function $F(\cdot)$; a terminal constraint is *not* used so that $X_f = \mathbb{R}^n$. It is generally necessary to have a terminal constraint (even if it can be omitted in the on-line computations if N is sufficiently large) if the system is nonlinear or if the system is linear and constrained but unstable. This leaves the linear, unconstrained case (Bitmead et al., 1990) and the linear, constrained, stable, case (Rawlings & Muske, 1993).

3.7.2.1. Linear, unconstrained systems. Here $f(x, u) = Ax + Bu$, and $\ell(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$ where $Q > 0$ and $R > 0$. Since the system is unconstrained, $\mathbb{X} = \mathbb{R}^n$ and $\mathbb{U} = \mathbb{R}^m$, so that A1–A3 are trivially satisfied. Let $\kappa_f(x) := K_f x$ stabilize the system ((A, B) is assumed stabilizable), and let $P_f > 0$ satisfy the Lyapunov equation

$$A_f^T P A_f + Q_f = 0, \quad A_f := A + B K_f,$$

$$Q_f := Q + K_f^T R K_f$$

Then $F(x) := (1/2)x^T P_f x$ satisfies A4 (with equality) and, with these ingredients, viz. $F(x) := (1/2)x^T P_f x, X_f = \mathbb{R}^n$ (no terminal constraint) and $\kappa_f(x) = K_f x$, the closed-loop system is asymptotically (exponentially) stable with a region of attraction \mathbb{R}^n .

3.7.2.2. Linear, constrained, stable systems. The data are the same as in the first case except that, in addition, the system is stable and constrained. The control constraint set is convex and compact, but there are no state or terminal constraints so $\mathbb{X} = X_f = \mathbb{R}^n$. It follows from A2 that $\kappa_f(\cdot)$, if linear, must satisfy $\kappa_f(x) \equiv 0$ which is, of course, the choice made in Rawlings and Muske (1993). With this choice, A1–A3 are trivially satisfied. Let $P_f > 0$ satisfy the Lyapunov equation

$$A^T P A + Q = 0.$$

Then $F(x) := (1/2)x^T P_f x$ satisfies A4 with equality. It follows that, with the ingredients $F(x) := (1/2)x^T P_f x, X_f = \mathbb{R}^n$ (no terminal constraint) and $\kappa_f(x) \equiv 0$, the closed-loop system is asymptotically (exponentially) stable with a region of attraction \mathbb{R}^n .

3.7.3. Terminal constraint set

This type of model predictive control, for nonlinear, constrained systems, employs a terminal constraint

$x(N) \in X_f$ but no terminal cost ($F(x) \equiv 0$). It was introduced in Michalska and Mayne (1993), where a variable horizon N was employed to control constrained, nonlinear continuous-time systems. Fixed horizon versions for constrained, nonlinear, discrete-time systems, discussed here, are analysed in Chisci et al. (1996) and Scokaert et al. (1999). Model predictive control is employed to drive the state to X_f in finite time; inside X_f , a local stabilizing controller $\kappa_f(x) = K_f x$ is employed which motivates the name *dual-mode MPC* that is sometimes employed. In the fixed horizon version, it makes more sense to employ a terminal cost (such versions of model predictive control are discussed next), so our discussion here is brief.

The purpose of the model predictive controller is to steer any state outside X_f to X_f so the controller is similar to terminal equality constraint MPC except that $\{0\}$ is replaced by X_f . A local controller $\kappa_f(\cdot)$ (usually, but not necessarily linear) and set X_f is chosen to satisfy A1–A3 and to steer any initial state $x \in X_f$ to the origin exponentially fast. This is a standard problem requiring stabilizability of the linearized system $x^+ = Ax + Bu$ and sufficient smoothness of $f(\cdot)$. To satisfy A4 it is necessary that $\ell(x, \kappa_f(x)) = 0$ in X_f . A suitable choice is

$$\ell(x, u) := \alpha(x) \bar{\ell}(x, u),$$

where $\bar{\ell}(x, u) \geq c(|(x, u)|)^2$ (e.g. $\bar{\ell}(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$ where $Q > 0, R > 0$) and $\alpha(\cdot)$ is the characteristic function of X_f^c , i.e. $\alpha(x) := 0, x \in X_f$ and $\alpha(x) := 1, x \in X_f^c$. Where X_f^c is the complement of X_f . Then, A4 is satisfied with equality. Since $\ell(x, u) \geq c_1 > 0$ for all $x \in X_f^c$, it follows from (3.1) that $x \in X_N$ is steered by $\kappa_N(\cdot)$ to X_f in finite time. The closed-loop system is exponentially stable with a domain of attraction X_N .

3.7.4. Terminal cost and constraint set

This variant of model predictive control employs both a terminal cost $F(\cdot)$ and a terminal constraint $x(N) \in X_f$ in the optimal control problem $\mathcal{P}_N(x)$ and is the version attracting most attention in current research literature. It has superior performance when compared with zero state and terminal constraint set MPC and can handle a much wider range of problems than terminal cost MPC. We include in this category those variants of MPC (e.g. Parisini & Zoppoli, 1995; Jadbabaie et al., 1999) where a terminal constraint set is *implicit* in $\mathcal{P}_N(x)$ but not required in the on-line version because parameters (particularly the horizon N) are chosen to ensure automatic satisfaction of the terminal constraint.

3.7.4.1. Linear, constrained systems. Ideally, the terminal cost $F(\cdot)$ should be chosen to be $V_\infty^0(\cdot)$ in which case $V_N^0(\cdot) = V_\infty^0(\cdot)$ and the virtues of infinite horizon optimal control are obtained. This is generally impossible. However, if the system is linear and constrained (Sznaier & Damborg, 1987; Chmielewski

& Manousiouthakis, 1996; Scokaert & Rawlings, 1998), $F(x)$ can be chosen to be equal to $V_{uc}^0(x) = (1/2)x^T P_f x$, the value function for the *unconstrained* infinite horizon optimal control problem $\mathcal{P}_{\infty}^{uc}(x)$ (with $\ell(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$, $Q > 0$, $R > 0$), $\kappa_f(x) = K_f x$ the optimal controller for $\mathcal{P}_{\infty}^{uc}(x)$ and X_f the output admissible set for $x^+ = A_f x$, $A_f := A + BK_f$ thus satisfying A1–A4 (A4 is satisfied with equality). With these ingredients, $F(x) = V_{uc}^0(x) = V_{\infty}^0(x)$ for all $x \in X_f$. The closed-loop system is exponentially stable with a domain of attraction X_N . The authors of these papers advocate choosing the horizon N on-line by increasing N until the optimal terminal state obtained by solving problem $\mathcal{P}_N(x)$ without the terminal constraint actually satisfies $x^0(N; x) \in X_f$. Although N varies with x in this version, exponential stability follows using $V_{\infty}^0(\cdot)$ as a Lyapunov function. The essential role of X_f is to permit replacement of $V_{\infty}^0(x)$ (the ideal choice for $F(x)$) by $V_{uc}^0(x) = (1/2)x^T P_f x$.

3.7.4.2. Nonlinear, unconstrained systems. This problem has been addressed, for example, by Parisini and Zoppoli (1995), Magni and Sepulchre (1997) and Jadbabaie et al. (1999); the literature on the nonlinear, constrained case, discussed next, is also relevant. Parisini and Zoppoli (1995) choose, as in Michalska and Mayne (1993), $\kappa_f(x) = K_f x$ to be a locally stabilizing controller and $x^T P_f x$ to be a Lyapunov function for the linearized system $x^+ = Ax + B\kappa_f(x)$. By choosing X_f to be a sufficiently small level set of $x \mapsto x^T P_f x$, A1 and A2 are satisfied and $x^T P_f x$ is also a Lyapunov function with a domain of attraction X_f for the nonlinear system $x^+ = f(x, \kappa_f(x))$. Parisini and Zoppoli (1995) now set $F(x) = ax^T P_f x$ so that A3 and A4 are satisfied and establish the existence of an integer N , a constant $a > 0$ and a level set of $V_N^0(\cdot)$ such that the terminal constraint $x(N) \in X_f$ may be omitted from the version of $\mathcal{P}_N(x)$ solved on-line for all states in this level set; the resultant closed-loop system $x^+ = f(x, \kappa_N(x))$ is asymptotically stable with a domain of attraction that is this (un-quantified) level set. Jadbabaie et al. (1999) give a definitive and illuminating discussion (in the context of continuous-time systems) of the unconstrained nonlinear problem. They show that the essential requirement for stability is that $F(\cdot)$ be a control Lyapunov function in the neighbourhood of the origin; if this is the case, there exists a constraint set $X_f = \{x \mid F(x) \leq r\}$ and a local control law $\kappa_f(\cdot)$ such that the triple $(F(\cdot), X_f, \kappa_f(\cdot))$ satisfies A1–A4. They omit the terminal constraint from $\mathcal{P}_N(x)$ that is, therefore, an unconstrained optimal control problem and establish, *inter alia*, the following useful (and surprising) results: (i) if $x \in X_f$, then $x^0(N; x) \in X_f$ and $V_N^0(x) \leq F(x)$ (see the appendix) and (ii) if $x \in \Gamma_r^N := \{x' \mid V_N^0(x') \leq r\}$ then $x^0(N; x) \in X_f = \{x' \mid F(x') \leq r\}$ so that the terminal constraint may be omitted from $\mathcal{P}_N(x)$ for all $x \in \Gamma_r^N$. This set is positively invariant under model predictive control so that the resultant closed-loop sys-

tem (with the terminal constraint omitted from $\mathcal{P}_N(x)$) is exponentially stable with a domain of attraction Γ_r^N . They also show that for any $s > 0$ there exists a horizon N_s such that $x^0(N_s; x) \in X_f$ for every initial state $x \in \Gamma_s^{\infty} := \{x' \mid V_{\infty}^0(x') \leq s\}$, thus permitting the region of attraction to be enlarged.

3.7.4.3. Nonlinear, constrained systems. If the system is nonlinear and constrained, features of model predictive control for both the above classes of problems appear (Chen & Allgöwer, 1998b; De Nicolao et al., 1996c; De Nicolao, Magnani, Magni & Scattolini, 1999a; De Nicolao, Magni & Scattolini, 1998). Thus, both Chen and Allgöwer (1998b) and De Nicolao et al. (1996c) (modulo the fact that continuous-time systems are considered in the first paper) choose $\kappa_f(x) = K_f x$ to stabilize the linearized system $x^+ = Ax + Bu$ and choose X_f to satisfy the set constraints $X_f \subset \mathbb{X}$ and $\kappa_f(X_f) \subset \mathbb{U}$; these choices satisfy A1 and A2. Then Chen and Allgöwer (1998b) *in effect* choose X_f to be a level set of $F(\cdot)$ where $F(\cdot) = (1/2)x^T P x$ is a control Lyapunov function for the linearized system $x^+ = Ax + Bu$ satisfying the Lyapunov equation

$$F(Ax + B\kappa_f(x)) - F(x) + \bar{\ell}(x, \kappa_f(x)) = 0$$

for all x where $\bar{\ell}(x, u) = \beta \ell(x, u)$, or $\bar{\ell}(x, u) = (\beta/2)(|x|_Q^2 + |u|_R^2)$ when $\ell(x, u) = (1/2)(|x|_Q^2 + |u|_R^2)$ and $\beta \in (1, \infty)$ (Chen and Allgöwer achieve this by replacing A by $A + \rho I$ in the continuous-time Lyapunov equation for the system $x^+ = Ax + Bu$). Replacing $\ell(\cdot)$ by $\bar{\ell}(\cdot) > \ell(\cdot)$ provides sufficient margin to ensure that A3 and A4 are satisfied when X_f is a sufficiently small level set of $F(\cdot)$. De Nicolao et al. (1996c), on the other hand, choose $F(\cdot)$ to be the infinite horizon cost incurred by the nonlinear system $x^+ = f(x, \kappa_f(x))$ and X_f to be a positively invariant region of attraction (for $x^+ = f(x, \kappa_f(x))$). This choice ensures satisfaction of A3 and A4. In both variants, A1–A4 are satisfied and closed-loop asymptotic stability results. The variant due to De Nicolao et al. (1996c) is conceptual since neither $F(\cdot)$ nor X_f can be precisely computed; an implementable version is given in De Nicolao, Magnani, Magni and Scattolini (1999a) in which the infinite horizon cost terminal cost $F(\cdot)$ defined above is approximated by a cost over the finite horizon M . Because of the approximation error, $F(\cdot)$, which now denotes the finite horizon terminal cost, no longer satisfies A4. To overcome this an X_f and $\kappa_f(\cdot)$, similar to those in Chen and Allgöwer (1998b), are chosen so that $\kappa_f(\cdot)$ is exponentially stabilizing in X_f . The horizon M (which depends on the current state x) is chosen sufficiently large to ensure that the approximation error is less than $\rho|x|_Q^2$ so that

$$[V_N^* + (1 - \rho)\ell](x, \kappa_N(x)) \leq 0$$

for all $x \in X_f$ where $\rho \in (0, 1)$ which is sufficient for closed-loop stability.

The paper (Alamir & Bornard, 1995) is, in some respects, similar; a finite control horizon N and cost (prediction) horizon M , $M > N$, are employed. This is equivalent to having a terminal cost of $F(x) = \sum_{k=N}^M \ell(x(k), u(k))$, but no terminal constraint. The omission of the latter appears to be due to the assumption that $\mathcal{P}_N^M(x)$ has a solution for all $x \in \mathbb{X}$, all $M \geq N$ making it unnecessary to establish the existence of a feasible solution at x^+ given that one exists at x . Alamir and Bornard (1995) establish the existence of a M_0 such that the closed-loop system is asymptotically stable, with a region of attraction \mathbb{X} , for all $M \geq M_0$.

3.8. Other forms of model predictive control

3.8.1. Variable horizon model predictive control

This was proposed in Michalska and Mayne (1993) in the context of dual-mode model predictive control and developed (Michalska, 1997); it has also been employed in contractive model predictive control (see the next section). In the dual-mode version, a terminal set X_f is chosen, and the optimal control problem solved on-line is

$$\mathcal{P}(x): \quad V^0(x) = \min_{(\mathbf{u}, N)} \{V_N(x, \mathbf{u}) \mid \mathbf{u} \in \mathcal{U}_N(x), \\ N \in \{1, \dots, N_{\max}\}\},$$

where $V_N(\cdot)$ is defined as before (1.10). The implicit model predictive control law is $\kappa(\cdot): X_f^c \rightarrow \mathbb{U}$. Because of the variable horizon, the value function automatically satisfies

$$[V^0 + \ell](x, \kappa(x)) \leq 0$$

(A1–A4 are not required). With our assumptions on $\ell(\cdot)$, the controller $\kappa(\cdot)$ steers any $x \in X_{N_{\max}} \setminus X_f$ to X_f in finite time. Inside X_f a local stabilizing controller $\kappa_f(\cdot)$ is employed; assumptions A1–A4 ensure its existence.

3.8.2. Contractive model predictive control

Contractive model predictive control, developed in a series of papers (Polak & Yang, 1993a,b; Yang & Polak, 1993) and Morari and De Oliveira (1998), is Lyapunov based but avoids the usual difficulties in obtaining a Lyapunov function by choosing, a priori, a positive-definite function $M(\cdot)$ and using on-line optimization, at event (x, k) , to ensure that $M(x^0(N^0(x); x)) \leq \beta M(x)$, $N^0(x) \geq 1$ where $\beta \in (0, 1)$. Usually, $M(x) = (1/2)x^T P x$ where $x \mapsto x^T P x$ is a control Lyapunov function for the linearized system (linearized at the origin). The variable horizon optimal open-loop control problem solved on-line at event (x, k) , is

$$\mathcal{P}^{\text{ctr}}(x): \quad V^0(x) = \min_{(\mathbf{u}, N)} \{V^*(x, \mathbf{u}, N) \mid \mathbf{u} \in \mathcal{U}_N(x), \\ N \in \{1, \dots, N_{\max}\}\},$$

where, as before, $\mathcal{U}_N(x)$ is the set of feasible (admissible) control sequences satisfying, for the given initial state x , the control, state and terminal constraints, except now that the terminal constraint $X_f(x)$ depends on x and is defined by

$$X_f(x) := \{x' \mid M(x') \leq \beta M(x)\}.$$

Solving $\mathcal{P}^{\text{ctr}}(x)$ yields the optimizing sequence $\mathbf{u}^0(x) = \{u^0(i; x)\}$ and the optimal horizon $N^0(x)$. In the original form of contractive MPC, the whole sequence of controls $\mathbf{u}^0(x) = \{u^0(0; x), u(k+1) = u^0(1; x), \dots, u(k+N^0(x)-1) = u^0(N^0(x)-1; x)\}$ is applied to the plant *open-loop* so that $u(k) = u^0(0; x)$, $u(k+1) = u^0(1; x), \dots$ and, in the absence of noise and model error, $M(x(k+N_x^0; x)) \leq \beta M(x(k))$ for all k where $x(k)$ is the plant state at time k . The procedure is repeated at time $k+N^0(x)$. If $\{t_j\}$ denotes the set of times when $\mathcal{P}^{\text{ctr}}(\cdot)$ is solved, it follows that, for all j , $M(x(t_j)) \leq \beta^j M(x(0))$. Exponential stability follows *provided* $\mathcal{P}^{\text{ctr}}(\cdot)$ can be solved at each event $(x(k), k)$ encountered, i.e. provided $x(k) \in X_{N_{\max}}$ for all k . This is not necessarily the case; the set of states in which $\mathcal{P}^{\text{ctr}}(x)$ has a feasible solution is not necessarily positively invariant under the model predictive control law. However, a sufficiently small level set of $V^0(\cdot)$ is positively invariant, and contractive model predictive control is asymptotically stabilizing with a region of attraction that includes this set. Morari and De Oliveira (1998) show how further optimization may be performed at times between the times $\{t_j\}$ when $\mathcal{P}^{\text{ctr}}(\cdot)$ is solved, thus providing extra feedback.

3.8.3. Stability enforced model predictive control

Stability of model predictive control systems is usually established using the value function as a Lyapunov function. A few authors have proposed the use of an additional Lyapunov function to enforce stability, allowing performance to be the main objective of model predictive control. Perhaps the earliest proposal for this version of model predictive control is due to Sznajder and Damborg (1990), who employ an auxiliary Lyapunov function $v(x)$ whose level set $\{x \mid v(x) \leq 1\}$ is the state constraint set \mathbb{X} . An extra stability constraint $v(f(x, u)) < v(x)$ is added to the optimal control problem $\mathcal{P}_N(x)$. Bemporad (1998a), in the context of constrained, linear, discrete-time systems, proposes the use of an auxiliary quadratic Lyapunov function; the descent property of the Lyapunov function is similarly enforced by an additional stability constraint in the optimal control problem $\mathcal{P}_N(x)$. The Lyapunov function is easily obtained but is not necessarily valid everywhere in X_N . Primbs, Nevistic and Doyle (1998), in the context of unconstrained nonlinear continuous-time systems, assume knowledge of a global control Lyapunov function $J(\cdot)$ from which they derive a globally stabilizing nonlinear controller $u = h_\sigma(x)$ satisfying $\dot{J}(x, h(x)) \leq -\sigma(x)$ where $\sigma(\cdot)$ is positive definite. The constraint $\dot{J}(x, u) \leq -\sigma(x)$ and a terminal

constraint that replaces X_f are added to the optimal control problem to enforce stability. It is difficult, of course, to obtain a suitable $J(\cdot)$.

3.8.4. System transformation

Linearization theory may, in some applications, be employed to transform the original nonlinear system, using state and feedback control transformations, into a linear system. Model predictive control may be applied to the transformed system (De Oliveira, Nevistic & Morari, 1995; Kurtz & Henson, 1997). The optimal control problem is not, however, transformed into a convex problem, because the transformed control and state constraint sets and the transformed cost are no longer necessarily convex. Keerthi (1986) and Rossiter, Kouvaritakis and Rice (1998) employ linear transformation ($x^+ = Ax + Bu$ is replaced by $x^+ = (A + BK)x + Bv$ where $v := u - Kx$ is the re-parameterized control) to improve conditioning of the optimal control problem $\mathcal{P}_N(x)$ solved on-line.

3.9. Suboptimal model predictive control

When the system is linear, the cost quadratic and the constraint sets are polyhedral, the optimal open-loop control problem reduces to a quadratic programme (X_f in $\mathcal{P}_N(x)$ may be replaced, if necessary, by a polytopic inner approximation) for which efficient programs exist. These yield a global solution to the optimal open-loop control problem, so that the model predictive controllers, described above, may be implemented. However, when the system is nonlinear, the optimal open-loop control problem is non-convex, and conventional nonlinear programming algorithms can only be relied on to give local, rather than global, solutions. The question then arises whether the desirable properties of model predictive control are lost if global solutions to the optimal open-loop control problem are not computable. The question is important because it is unlikely that global optimization algorithms, despite the progress being made in this field, will be employed for large problems. Indeed, it generically requires an infinite amount of computation merely to verify that a point x is globally optimal. The situation is vastly better for finding feasible solutions; the measure of the set of feasible solutions is generically positive, in contrast to zero for globally optimal solutions (if unique), and verification of the feasibility of a point x is simple. So, model predictive controllers that require feasibility rather than optimality have a much better prospect of being implemented when the system is nonlinear.

The problem is addressed in Michalska and Mayne (1993) and Mayne (1995) where it was shown, for continuous-time, variable horizon, dual-mode model predictive control, that it is possible to ensure stability by using a feasible solution to the open-loop control problem that is merely better, in a well defined sense, than the preced-

ing solution; optimality is *not* required, merely feasibility. Chisci et al. (1996) extend this strategy to fixed horizon model predictive control for discrete-time systems. A more general analysis (Scokaert et al., 1999) establishes that, under mild conditions, feasibility rather than optimality suffices for stability, so that many existing model predictive controllers, for which stability has been established under the assumption that global solutions to the optimal open-loop control problem are employed, may, possibly with minor modifications, retain their stabilizing properties when global minimization is not possible. The purpose of this result is not to advocate suboptimality (optimality is preferred when possible), but to permit satisfactory control when achieving optimality is impractical. Many strategies are possible. For example, the controller may attempt to find an optimal solution to the optimal control problem and cease when a permitted time limit is reached. Or it may solve a simpler version of the optimal control problem, for example one in which only the first few control actions in the sequence are optimized (Zheng, 1997).

3.10. Conclusion

The surprising, and happy, discovery of this survey is the rapidly emerging consensus reached by many researchers on the essential ingredients (terminal cost, $F(\cdot)$, terminal constraint set X_f and local stabilizing controller $\kappa_f(\cdot)$) that ensure closed-loop stability when model predictive control is employed. The first two ingredients appear explicitly in the optimal control problem (although automatic satisfaction of the terminal constraint may be attained by prior choice of a sufficiently large horizon). The local controller $\kappa_f(\cdot)$ is merely implicit, but is required to prove stability (and to provide a ‘hot’ start for the optimal control problem $\mathcal{P}_N(x)$). We have distilled from this consensus a set of four conditions (or axioms) A1–A4 that are sufficient for closed-loop stability (modulo further modest assumptions). This provides a useful unification of existing work, in that most variants of model predictive control differ only in their choice of these ingredients.

It appears desirable, for many reasons, to have a terminal cost $F(\cdot)$ which is as close to the value function $V_\infty^0(\cdot)$ as possible. This choice, if exactly achieved, gives the known benefits of infinite horizon optimal control. One benefit, not discussed above, that seems important in the context of model predictive control, is that this choice ensures the closed-loop trajectory of the plant is exactly that predicted by the solution of the optimal control problem $\mathcal{P}_N(x)$ so that the effect of the cost parameters on (nominal) closed-loop performance is transparent. This correspondence of predicted and actual trajectories is not achieved in all versions of model predictive control; zero terminal state model predictive control can, if the horizon N is less than the plant’s settling

time, result in a plant trajectory that differs considerably from its prediction. Since $V_{\infty}^0(\cdot)$ is not known (except if the system is linear and unconstrained and the cost quadratic), a good approximation in a neighbourhood X_f of the target state is required. When the system is linear and constrained, the set X_f can be chosen so that $V_{\infty}^0(\cdot)$ is exactly known in it; hence, if N is chosen appropriately, the closed-loop trajectory of the plant is exactly that predicted by the solution of the optimal control problem $\mathcal{P}_N(x)$. When the system is nonlinear, constrained or unconstrained, $F(\cdot)$ and X_f may be chosen so that $F(x)$ is a good approximation to $V_{\infty}^0(x)$ in X_f . A useful result of satisfaction of A1–A4 is that the finite horizon optimal control problem $\mathcal{P}_N(x)$ is equivalent to a modified infinite horizon optimal problem (in which $\ell(\cdot)$ is replaced by $\bar{\ell}(\cdot) > \ell(\cdot)$). So the known benefits of infinite horizon optimal control are achieved (if A1–A4 are satisfied) even though the horizon is finite.

4. Robustness

4.1. Introduction

The introduction of uncertainty in the system description raises the question of robustness, i.e. the maintenance of certain properties such as stability and performance in the presence of uncertainty. Most studies on robustness consider unconstrained systems; if a Lyapunov function for the nominal closed-loop system maintains its descent property if the disturbance (uncertainty) is sufficiently small, then stability is maintained in the presence of uncertainty. However, when constraints on states and controls are present, it is necessary to ensure, in addition, that disturbances do not cause transgression of the constraints; this adds an extra level of complexity.

The earliest analyses of robustness of model predictive control employed impulse response models. Richalet et al. (1978) investigate robustness in the face of gain mismatch. In later papers, control is determined by solving a min–max optimal control problem where the ‘adversary’ represents uncertainty in the impulse response. The complexity of this problem increases exponentially with horizon length although Allwright (1993) shows how this complexity may be substantially reduced. For further discussion of robustness of model predictive control using impulse response models, see Zheng and Morari (1993), Genceli and Nikolaou (1993), and De Nicolao, Magni and Scattolini (1996b). It is relatively simple to ensure robust stability for linear or nonlinear systems that have finite memory (finite impulse response models, finite Volterra models, etc.) since a stability constraint (e.g. $x(N) = 0$) is simply implementable (by choosing N and k_1 appropriately and imposing the constraint, in the optimal control problem, that the control is zero for all $k \geq k_1$).

There are several approaches to the study of robustness. The first is concerned with the robustness of closed-loop systems, designed using the nominal system (i.e. neglecting uncertainty). The second attempts to achieve robustness in the context of conventional model predictive control by consideration of all possible realizations on the uncertainty (min–max open-loop model predictive control). A defect of model predictive control of uncertain systems, not yet widely appreciated, is the open-loop nature of the optimal control problem; the third approach addresses this by introducing feedback in the min–max optimal control problem solved on-line. We briefly address each approach below.

4.2. Modelling uncertainty

We shall suppose the uncertain system is described by

$$x^+ = f(x, u, w),$$

$$z = h(x).$$

The state x and control u satisfy the same constraints as before, and the adversary w satisfies $w(k) \in W(x(k), u(k))$ for all k where, for each (x, u) , $W(x, u)$ is closed (perhaps compact) and contains the origin in its interior. Because $f(\cdot)$ now depends on w , we define

$$* \phi(x, u, w) := \phi(f(x, u, w)) - \phi(x),$$

where $\phi(x)$ is any function. Let $w(\cdot)$ or $\mathbf{w} := \{w(0), w(1), \dots, w(N-1)\}$ denote a disturbance sequence and $x^{u, \mathbf{w}}(\cdot; x)$ the state trajectory (sequence) resulting from an initial state x at time 0, and control and disturbance sequences \mathbf{u} and \mathbf{w} , respectively. Let $\mathcal{F}(x, u) := f(x, u, W(x, u))$; then $\mathcal{F}(\cdot)$ maps points in $\mathbb{X} \times \mathbb{U}$ to subsets of \mathbb{R}^n and $x^+ \in \mathcal{F}(x, u)$ is an alternative definition of the system. In some situations the uncertainty (e.g. state estimation error) is time varying in which case uncertainty may be better modelled by $w(k) \in W_k$ where W_k varies appropriately with time k .

4.3. Inherent robustness

By inherent robustness we mean robustness of the closed-loop system using model predictive control obtained *ignoring* uncertainty. This has been investigated in De Nicolao et al. (1996a) and Magni and Sepulchre (1997). We discuss the latter because the continuous-time case is simpler. The unconstrained system (ignoring uncertainty) is

$$\dot{x} = f(x) + g(x)u$$

and the cost (again ignoring uncertainty)

$$V(x, u(\cdot)) = \int_0^T \ell(x(s), u(s)) ds + F(x(T)),$$

where $x(s) = x^u(s; x, 0)$, $\ell(x, u) = (1/2)|u|_R^2 + q(x)$ and $q(\cdot)$ is positive definite. There are no constraints on x and u but there is a terminal constraint $x(T) \in X_f$. The optimal control $\kappa_T(x)$ minimizes $\ell(x, u) + \langle \nabla V_T^0(x), f(x) + g(x)u \rangle$ so that

$$\kappa_T(x) = -R^{-1}g(x)^T \nabla V_T^0(x).$$

If B1–B4 are satisfied

$$\dot{V}_T^0(x) + \bar{q}(x) + (1/2)|\kappa_T(x)|_R^2 = 0,$$

where $\bar{q}(x) = q(x) + (\partial/\partial t)V^0(x, 0) \geq q(x)$; thus $\kappa_T(\cdot)$ is optimal for a (modified) infinite horizon optimal control problem. Standard calculations now show that $\dot{V}_T^0(x) + \bar{q}(x) \leq 0$ if the system is replaced by

$$\dot{x} = f(x) + g(x)\phi(u),$$

where $\phi(\cdot)$ is a nonlinearity in the sector $(1/2, \infty)$. Hence, the closed-loop system is asymptotically stable with a region of attraction $\{x | V_T^0(x) \leq c\} \subset X_N$; the region of attraction may be considerably smaller than X_N .

4.4. Conditions on $F(\cdot)$, X_f and $\kappa_f(\cdot)$

Subsequent versions of robust model predictive control that consider *all* realizations of the disturbance sequence \mathbf{w} in the optimal control problem require strengthened assumptions. The ingredients $F(\cdot)$, X_f and $\kappa_f(\cdot)$ are therefore assumed, in the sequel, to satisfy *robust* versions of A1–A4, viz.,

A1: $X_f \subset \mathbb{X}$, X_f closed, $0 \in X_f$.

A2: $\kappa_f(x) \in \mathbb{U}$, $\forall x \in X_f$.

A3a: $f(x, \kappa_f(x), w) \in X_f$, $\forall x \in X_f$, $\forall w \in W(x, \kappa_f(x))$.

A4a: $[F + \ell]^*(x, \kappa_f(x), w) \leq 0$, $\forall x \in X_f$, $\forall w \in W(x, \kappa_f(x))$.

There exist such a triple if $F(\cdot)$ is a robust control Lyapunov function in a neighbourhood of the origin. These assumptions ensure

$$[V_N^* + \ell](x, \kappa_N(x), w) \leq 0$$

(or, equivalently, the monotonicity property) for all x in an appropriate set and all $w \in W(x, \kappa_N(x))$ and, hence, asymptotic or exponential stability.

4.5. Open-loop min–max model predictive control

Deterministic model predictive control, discussed in Section 3, has the property that X_N is positively invariant for the closed-loop system $x^+ = f(x, \kappa_N(x))$; if $x \in X_N$, then $x^+ \in X_{N-1} \subset X_N$. This property is lost when uncertainty is present, as illustrated in Section 4.3. To recover this property, it is necessary to consider all possible realizations of $x^+ \in \mathcal{F}(x, u)$ in the optimal control problem and ensure each realization satisfies the state, control and terminal constraints (Michalska & Mayne, 1993;

Mayne, 1995; Chen, Scherer & Allgöwer, 1997; Magni, Nijmeijer & van der Schaft, 1999b; Magni, De Nicolao, Scattolini & Allgöwer, 1999a). The cost of an individual realization is defined by

$$J(x, \mathbf{u}, \mathbf{w}) := \sum_{i=0}^{N-1} \ell(x(s), u(s)) + F(x(N)),$$

where, now, $x(s) = x^{u, w}(s; x, 0)$ and the cost is

$$V_N(x, \mathbf{u}) := \max\{J(x, \mathbf{u}, \mathbf{w}) | \mathbf{w} \in \mathcal{W}_N(x, \mathbf{u})\},$$

where $\mathcal{W}_N(x, \mathbf{u})$ is the set of admissible disturbance sequences. Other choices are sometimes made; for example $V(x, \mathbf{u})$ may be the cost $J(x, \mathbf{u}, \mathbf{0})$ of the nominal system where $\mathbf{0}$ is the zero sequence. An interesting variant, in the context of model predictive control of stable linear uncertain systems, is proposed in Badgwell (1997); here the optimal control problem incorporates a robust stability constraint, namely that the control reduces the cost associated with *each* possible realization (assumed finite in number) of the linear system. This is stronger than merely a reduction in the maximum cost. In some cases (H_∞ model predictive control), $\ell(\cdot)$ is a function of (x, u, w) . Let $\mathcal{U}_N^{\text{ol}}(x)$ now denote the set of admissible control sequences \mathbf{u} satisfying the state, control and terminal constraints for *every* admissible disturbance sequence \mathbf{w} when the initial state is x . Clearly $\mathcal{U}_N^{\text{ol}}(x) \subset \mathcal{U}_N(x)$. For all $i \geq 0$, let X_i^{ol} denote the set of states x such that $\mathcal{U}_N^{\text{ol}} \neq \emptyset$; $X_i^{\text{ol}} \subset X_i$ is the set of states that can be robustly steered to X_f in i steps or less by an admissible control sequence \mathbf{u} . The open-loop optimal control problem is

$$\mathcal{P}_N^{\text{ol}}(x): V_N^{\text{ol}}(x) = \min\{V_N(x, \mathbf{u}) | \mathbf{u} \in \mathcal{U}_N^{\text{ol}}(x)\}.$$

The solution $\mathbf{u}^{\text{ol}}(x)$ of $\mathcal{P}_N^{\text{ol}}(x)$ yields the implicit min–max model predictive control law:

$$\kappa_N^{\text{ol}}(x) := u^{\text{ol}}(0; x)$$

as before. Corresponding to $\mathbf{u}^{\text{ol}}(x)$ is a ‘bundle’ of optimal state trajectories $\{x^{\text{ol}}(x, \mathbf{w})\}$, one for each admissible \mathbf{w} , where

$$x^{\text{ol}}(x, \mathbf{w}) = \{x^{\text{ol}}(0; x, \mathbf{w}), x^{\text{ol}}(1; x, \mathbf{w}), \dots, x^{\text{ol}}(N; x, \mathbf{w})\}.$$

By the definition of $\mathcal{P}_N^{\text{ol}}(x)$, $x^{\text{ol}}(N; x, \mathbf{w}) \in X_f$ for each admissible \mathbf{w} . The ingredients $(F(\cdot), X_f, \kappa_f(\cdot))$ are assumed to satisfy A1, A2, A3a and A4a. At this point a difficulty arises: suppose $x \in X_N^{\text{ol}}$ and an optimal (hence admissible) control sequence $\{u^{\text{ol}}(0; x), u^{\text{ol}}(1; x), \dots, u^{\text{ol}}(N-1; x)\}$ for $\mathcal{P}_N^{\text{ol}}(x)$ is known; this sequence steers every realization of the state (corresponding to different realizations of \mathbf{w}) to X_f in N steps or less so that $x^{\text{ol}}(N; x, \mathbf{w}) \in X_f$ for every $\mathbf{w} \in \mathcal{W}(x, \mathbf{u}^{\text{ol}}(x))$. Hence, the abbreviated control sequence $\{u^{\text{ol}}(1; x), \dots, u^{\text{ol}}(N-1; x)\}$ steers every $x^+ \in \mathcal{F}(x, \kappa_N(x))$ to X_f in $N-1$ steps or less so that $x^+ \in X_{N-1}^{\text{ol}}$. The difficulty is obtaining a feasible control *sequence*

$$\tilde{u}(x) = \{u^{\text{ol}}(1; x), \dots, u^{\text{ol}}(N-1; x), v\}$$

for $\mathcal{P}_N(x^+)$ where x^+ is any state in $\mathcal{F}(x, \kappa_N(x))$; the control action $v \in \mathbb{U}$ is required to satisfy $f(x^0(N; x, \mathbf{w}), v\omega_N) \in X_f$ for every $\mathbf{w} \in \mathcal{W}(x, \mathbf{u}^0(x))$. Condition A3a does not ensure the existence of a v with this property. Without a feasible sequence, an upper bound for $V_N^0(x^+)$ cannot be obtained. One way of avoiding this impasse is to replace the fixed horizon strategy by a variable horizon strategy (Michalska & Mayne, 1993) in which the decision variable is (\mathbf{u}, N) . Suppose $(\mathbf{u}^0(x), N^0(x))$ solves the resultant variable horizon optimal control problem $\mathcal{P}(x)$, where

$$\mathbf{u}^0(x) := \{u^0(0; x), u^0(1; x), \dots, u^0(N^0(x) - 1; x)\}.$$

Then $(\bar{\mathbf{u}}(x), N^0(x) - 1)$, where

$$\bar{\mathbf{u}}(x) := \{u^0(1; x), \dots, u^0(N^0(x) - 1; x)\}$$

is a feasible solution for the optimal control problem $\mathcal{P}(x^+)$ at any $x^+ \in \mathcal{F}(x, \kappa_N(x))$. With $V^0(\cdot)$ denoting the resultant value function and $\kappa^{ol}(\cdot)$ the resultant model predictive control law,

$$[V^0 + \ell](x, \kappa^{ol}(x), w) \leq 0$$

(or, equivalently, the monotonicity property) holds for all $x \in X_N^{ol} \setminus X_f$, all $w \in W(x, \kappa^{ol}(x))$. Inside X_f , a local robustly stabilizing control law $\kappa_f(\cdot)$ is used; conditions A1, A2, A3a and A4a ensure the existence of a suitable X_f and $\kappa_f(\cdot)$. With further modest assumptions, robust asymptotic (exponential) stability results with a domain of attraction X_N^{ol} .

4.6. Feedback model predictive control

4.6.1. Introduction

Although open-loop min-max model predictive control appears attractive, it may be very conservative due to the open-loop nature of $\mathcal{P}_N^{ol}(x)$; for given (x, \mathbf{u}) , the trajectories satisfying $x^+ \in \mathcal{F}(x)$ may diverge (consider the system $x^+ = x + w$, $w \in [-1, 1]$) causing X_N^{ol} to be small or even empty for reasonable N . This is unrealistic, because feedback prevents the trajectories from diverging excessively; the scenario generated in solving $\mathcal{P}_N^{ol}(x)$ does not model accurately the uncertain control problem because it ignores feedback by searching over open-loop control sequences \mathbf{u} in minimizing $V(x, \mathbf{u})$. For this reason *feedback* model predictive control was proposed in Mayne (1995, 1997), Kothare, Balakrishnan and Morari (1996), Lee and Yu (1997), Scokaert and Mayne (1998), De Nicolao, Magni and Scattolini (1999b), Magni, Nijmeijer and van der Schaft (1999b), and Magni, De Nicolao, Scattolini and Allgöwer (1999a). In feedback model predictive control, the decision variable \mathbf{u} , which is a sequence of control actions, is replaced by a *policy* π which is a sequence of control laws. More precisely

$$\pi := \{u(0), \kappa_1(\cdot), \dots, \kappa_{N-1}(\cdot)\}$$

where, for each i , $\kappa_i(\cdot): \mathbb{X} \rightarrow \mathbb{U}$ is a control law whereas $u(0)$ is a control action (since there is only one initial state). The cost is

$$V_N(x, \pi) := \max\{J(x, \pi, \mathbf{w}) \mid \mathbf{w} \in \mathcal{W}_N(x, \pi)\},$$

where

$$J(x, \pi, \mathbf{w}) := \sum_{i=0}^{N-1} \ell(x(s), u(s)) + F(x(N))$$

with $x(s) = x^{\pi, \mathbf{u}}(s; x)$, the state at time s resulting from an initial state x at time 0, a control policy π and a disturbance sequence \mathbf{w} ; $\mathcal{W}_N(x, \pi)$ is the set of admissible disturbance sequences given that control policy π is employed. Let $\Pi_N(x)$ denote the set of admissible control policies π satisfying the state, control and terminal constraints for every admissible disturbance when the initial state is x . For all $i \geq 0$, let X_i^{fb} denote the set of states x such that $\Pi_i \neq \emptyset$; $X_i^{\text{fb}} \subset X_i$ is the set of states that can be robustly steered to X_f in i steps or less by an admissible policy π . The feedback optimal control problem is

$$\mathcal{P}_N^{\text{fb}}(x): V_N^0(x) = \min\{V_N(x, \pi) \mid \pi \in \Pi_N(x)\}.$$

The solution, if it exists, is

$$\pi^0(x) = \{u^0(0; x), \kappa_1^0(\cdot; x), \dots, \kappa_{N-1}^0(\cdot; x)\}$$

of $\mathcal{P}_N^{\text{fb}}(x)$ yields an implicit feedback model predictive control law

$$\kappa_N^{\text{fb}}(x) := u^0(0; x).$$

Suppose assumptions A1, A2, A3a and A4a hold. Then

$$\tilde{\pi}(x, x^+) := \{\kappa_1^0(x^+; x), \kappa_2^0(\cdot; x), \dots, \kappa_{N-1}^0(\cdot; x), \kappa_f(\cdot)\}$$

is a feasible policy for $\mathcal{P}_N^{\text{fb}}(x^+)$ for all $x^+ \in \mathcal{F}(x, \kappa_N^{\text{fb}}(x))$, all $x \in X_N^{\text{fb}}$, X_N^{fb} is positively invariant for the system $x^+ \in \mathcal{F}(x, \kappa_N^{\text{fb}}(x))$ and

$$[V_N^0 + \ell](x, \kappa_N^{\text{fb}}(x), w) \leq 0$$

for all $x \in X_N^{\text{fb}}$, all $w \in W(x, \kappa_N^{\text{fb}}(x))$. With further modest assumptions, robust asymptotic (exponential) stability results with a domain of attraction X_N^{fb} . The results are similar to those for open-loop min-max model predictive control; the gain is that X_N^{fb} contains X_N^{ol} (and is possibly far larger).

The feedback version of model predictive control appears attractive but prohibitively complex. It is encouraging, therefore, to note two examples that have appeared in the recent literature.

4.6.2. Linear feedback MPC

An interesting proposal (Kothare et al., 1996) for robust model predictive control of linear constrained systems that recognizes the need for a feedback version of the optimal control problem solved on-line has recently

appeared. An abbreviated description follows. The system is

$$x^+ = Ax + Bu, \quad y = Cx,$$

where $\theta := (A, B)$ is only known to lie in a convex set Θ , the convex hull of $\{(A_i, B_i) \mid i \in I\}$. Let $F(x, P) := x^T P x$, $A_K := A + BK$ and $\ell(x, u) := x^T Q x + u^T R u$ where Q and R are positive definite. The control and state constraints are $\mathbb{U} := \{u \mid \|u\| \leq c_u\}$ and $x \in \mathbb{X} := \{x \mid \|Cx\| \leq c_x\}$. At each state x , the (feedback) optimal control problem $\mathcal{P}(x)$ of minimizing γ with respect to (γ, P, K) subject to the constraints

- (i) $P > 0$, $F(x, P) \leq \gamma$,
- (ii) $F(A_K x', P) - F(x', P) + \ell(x', Kx') \leq 0$, $\forall x' \in X(\gamma, P)$,
 $\forall (A, B) \in \Theta$,
- (iii) $|Kx'| \leq c_u$ and $|Cx'| \leq c_x$, $\forall x' \in X(\gamma, P)$,

where $X(\gamma, P) := \{x \mid F(x, P) \leq \gamma\}$ is solved yielding $(\gamma^0(x), P^0(x), K^0(x))$. This problem is shown in Kothare et al. (1996) to be equivalent to a linear matrix inequality. The control $u = K^0(x)x$ is applied to the plant, and the procedure repeated. The controller generates a sequence $\{x(k), P(k), \gamma(k)\}$ satisfying $x(k) \in X(\gamma(k), P(k))$ and $\gamma(k+1) \leq \gamma(k) - x(k)^T Q x(k)$ for all k and all admissible parameter sequences. It readily follows that $x(k) \rightarrow 0$ as $k \rightarrow \infty$. This is an interesting contribution, albeit conservative because the sets $W(k)$, being ellipsoids, are conservative inner approximations to X_∞ , the set of states steerable to the origin and because the feedback policy π is a sequence of identical feedback laws.

4.6.3. Min-max feedback MPC

The second example (Scokaert & Mayne, 1998) addresses model predictive control of the uncertain system

$$x^+ = Ax + Bu + w,$$

where the state, control and disturbance satisfy the constraints $x \in \mathbb{X}$, $u \in \mathbb{U}$ and $w \in W$, respectively. These sets are convex, compact and contain the origin in their interiors; W is a fixed polytope ($W = \text{co}\{v_i \mid i \in I\}$) so the class of permissible disturbances is $\mathcal{W} = W^N$. Because the disturbance is merely bounded, the most that can be achieved is to steer the state to a ‘robust control invariant set X_f ’; a local controller $\kappa_f(\cdot)$ then maintains the state in X_f despite the disturbance. The optimal control problem $\mathcal{P}_N^{\text{fb}}(x)$ is of the terminal constraint variety with $F(x) \equiv 0$ and $\ell(\cdot)$ convex and zero in X_f . However, in solving $\mathcal{P}_N^{\text{fb}}(x)$, \mathcal{W} is replaced by $\mathcal{V} = V^N$ where $V = \{v_i \mid i \in I\}$ is the set of vertices of W and so has a finite number of elements $M = I^N - 1$. In this modified problem, instead of determining, at each $i \in \{0, 1, \dots, N-1\}$, a control law $\kappa^i(\cdot)$, an optimal control action $u_j^0(i; x)$, is determined for each of the I^i states $x_j^0(i; x)$ resulting from the different realizations of w so that the decision variable is

$$\pi = \{u(0), \{u_j(1) \mid j \in I\}, \dots, \{u_j(N-1) \mid j \in I^{N-1}\}\}.$$

Optimization yields $\pi^0(x)$, the optimal solution of $\mathcal{P}_N^{\text{fb}}(x)$ (with \mathcal{W} replaced by $\mathcal{V} = V^N$); the model predictive control law is $\kappa_N^{\text{fb}}(x) = u^0(0; x)$. An admissible $\tilde{\pi}^0(x, x^+)$ for $\mathcal{P}_N^{\text{fb}}(x^+)$ is then obtained by interpolation and incorporation of $\kappa_f(\cdot)$, yielding

$$[V_N^* + \ell](x, \kappa_N^{\text{fb}}(x), w) \leq 0$$

for all $x \in X_N^{\text{fb}}$, all $w \in W(x, \kappa_N^{\text{fb}}(x))$ and, hence, closed-loop stability. A defect of this approach is computational complexity; the number of vertices of \mathcal{W} is exponential in N .

4.7. H_∞ model predictive control

The earliest consideration of H_∞ model predictive control appears in Tadmor (1992) and Lall and Glover (1994), both papers being concerned with linear, time-varying, unconstrained, continuous-time systems. The time-varying nature of the problem, as in Kwon et al. (1983) justifies the use of model predictive control. Tadmor implicitly uses the terminal equality constraint $x(t+T) = 0$ (equivalently a terminal cost which is infinite except at the origin). Interestingly, Lall and Glover (1994) employ a (time-varying) quadratic terminal cost $F(x, t) = (1/2)x^T P(t)x$, a local time-varying controller $u = \kappa_f(x, t)$ and $X_f(t) = \mathbb{R}^n$ such that a (continuous-time) time-varying extension of A4a, viz.

$$[\dot{F} + \ell](x, \kappa_f(x), w, t) \leq 0$$

is satisfied for all $t > 0$, all $x \in \mathbb{R}^n$ and all w . These papers are also interesting in their use of feedback policy for the control in the optimal control problem solved on-line. Nonlinear H_∞ control, like its linear counterpart, ensures robustness by its choice of a cost function. However, determining the controller requires, in general, solution of a nonlinear Hamilton–Jacobi–Isaacs (HJI) equation. It is, therefore, not surprising that attempts have been made (Chen et al., 1997; De Nicolao, Magni & Scattolini, 1999b; Magni, Nijmeijer & van der Schaft, 1999b; Magni, De Nicolao, Scattolini & Allgöwer, 1999a) to avoid this complexity by using an appropriate version of model predictive control. The results are interesting, but implementation remains a difficulty. Following Chen et al. (1997), all utilize a terminal cost $F(\cdot)$ derived from H_∞ control of the linearized system and a terminal constraint set X_f . Chen et al. (1997) consider continuous-time, nonlinear systems and use an open-loop model predictive strategy, analogous to that described in Section 5.4. The remaining papers use the feedback model predictive control strategy described in Section 5.5, both in a continuous-time context (Magni, Nijmeijer & van der Schaft, 1999b) and a discrete-time context (De Nicolao, Magni & Scattolini, 1999b; Magni, De Nicolao, Scattolini & Allgöwer, 1999a). The underlying idea is

similar to that already discussed, the major change being the choice of the stage cost $\ell(\cdot)$. Consider the discrete-time system

$$x^+ = f(x, u, w), \quad z = (Hx, u),$$

with the cost of a particular realization, when policy π is employed, being

$$J(x, \pi, \mathbf{w}) := \sum_{i=0}^{N-1} \ell(z(i), w(i)) + F(x(N))$$

where $x(i) = x^{\pi, \mathbf{w}}(i; x, 0)$ and

$$\ell(z, w) := (\frac{1}{2})(|z|^2 - \gamma^2|w|^2).$$

The cost is defined by

$$V_N(x, \pi) := \max_w \{J(x, \pi, \mathbf{w}) \mid \mathbf{w} \in \mathcal{W}_N(x, \pi)\},$$

where $\mathcal{W}_N(x, \pi)$ is the set of admissible disturbances defined below. The optimal control problem is

$$\mathcal{P}_N(x): V_N^0(x) = \min\{V_N(x, \pi) \mid \pi \in \Pi_N(x)\},$$

where $\Pi_N(x)$ is the set of admissible policies. To simplify stability arguments we make the (very restrictive) assumption that $W(x, u)$ is defined by

$$W(x, u) := \{w \mid |w| \leq \delta|(Hx, u)|\}$$

and that $Q = H^T H > 0$. Hence

$$\begin{aligned} \ell(z, w) &= (1/2)(|z|^2 - \gamma^2|w|^2) \geq (1/2)((1 - \delta^2\gamma^2)|z|^2) \\ &= (\sigma/2)|z|^2 \geq (\mu/2)|x|^2 \end{aligned}$$

for all x, u and all $w \in W(x, u)$ where $\sigma > 0$ and $\mu > 0$ if $\delta < 1/\gamma$. This assumption may be relaxed (e.g. we may merely require that Δ is a dynamic system with ℓ_2 gain not exceeding δ and satisfaction of an observability condition) if more sophisticated stability arguments (van der Schaft, 1996) are employed.

To determine X_f and $F(\cdot)$, we consider a linear H_∞ problem for which the system is

$$x^+ = Ax + Bu + Gw,$$

where $A = f_x(0, 0)$, $B = f_u(0, 0)$ and $G = f_w(0, 0)$ and the cost of a realization is

$$\sum_{i=0}^{\infty} \ell'(x(i), u(i), w(i)),$$

where the stage cost

$$\ell'(x, u, w) = \ell((Hx, w), w) + (1/2)\rho|x|^2, \quad \rho > 1$$

is the same as that for the original nonlinear system except that Q is replaced by $Q + \rho I$ to provide a margin for approximation errors. If the data are appropriate (in particular if γ is large enough), the value function of the linear H_∞ problem is $F(x) = (1/2)x^T P x$, $P > 0$ and the

optimal control strategy is $u = \kappa_f(x)$. The value function $F(\cdot)$ satisfies

$$\begin{aligned} F(Ax + B\kappa_f(x) + Gw) - F(x + \ell((x, \kappa_f(x)), w) \\ + (1/2)\rho|x|^2) \leq 0 \end{aligned}$$

for all x , all admissible w . If $f(\cdot)$ is sufficiently smooth, there exists a level set of $F(\cdot)$, viz.

$$X_f := \{x \mid F(x) \leq c\}$$

such that

$$[F^* + \ell](x, \kappa_f(x), w) \leq 0$$

for all $x \in X_f$, all $w \in W(x, u)$. Under our assumptions, the triple $(F(\cdot), X_f, \kappa_f(\cdot))$ satisfies A1, A2, A3a and A4a so that

$$[V_N^{*0} + \ell](x, \kappa_N^{fb}(x), w) \leq 0$$

for all $x \in X_N^{fb}$, all $x^+ \in \mathcal{F}(x, \kappa_N^{fb}(x))$ and all $w \in W(x, u)$. Since $\ell(x, u, w) \geq (\mu/2)|x|^2$ for all admissible $w \in W(x, u)$, asymptotic (exponential) stability of the closed-loop system follows (under appropriate further assumptions). A further consequence of the last inequality is

$$\|z\|_2^2 \leq \gamma^2\|w\|_2^2 + V_N^0(x)$$

so that the closed-loop system ($u = \kappa_f^{fb}(x)$) has a L_2 gain (from w to z) not exceeding γ .

4.8. Conclusion

Research on robustness has clarified our understanding of the problem and the limitations inherent in conventional (open-loop) model predictive control. Perhaps these limitations should have been appreciated earlier; after all, it is not possible to define a linear quadratic optimal control problem that yields the H_∞ controller — a min-max optimal control problem must be solved where the minimization is over control strategies and not control sequences. The studies reported above define more clearly the problems that must be solved but, unlike research on stability, do not yet provide implementable solutions since the decision variable (π, \mathbf{w}) is infinite-dimensional. Feedback model predictive controllers must still be regarded as conceptual rather than practical. A tantalizing possibility for evading this impasse is presented by Gallestey and James (1999) in the context of *unconstrained* continuous-time systems of the form $\dot{x} = f(x) + g(x)u + h(x)w$. If a *global* control Lyapunov function $F(\cdot)$ is available for the terminal cost there is no need for a terminal constraint; the resultant unconstrained optimal control problem can then be formulated as a two-point boundary value problem solvable by shooting methods.

5. Other issues

5.1. Introduction

There remains a wealth of topics we have not discussed in this review. Our purpose here is to mention, rather than review, some important issues currently receiving attention from researchers.

5.2. Tracking

In the discussion above, the origin represents, with a suitable change of coordinates, any desired equilibrium state x_r such that $x_r \in \mathbb{X}$ and $u_r \in \mathbb{U}$ where u_r is the equilibrium control ($x_r = f(x_r, u_r)$). When the system being controlled is linear, and constant input and output disturbances are present, Muske and Rawlings (1993) utilize estimates of the disturbances to compute a (x_r, u_r) such that the equilibrium output y_r is as close as possible to the desired set-point r ($y_r = r$ if this is possible without transgressing constraints). For further reading on tracking constant reference signals see Lee, Morari and García (1994), Lee and Cooley (1997), Rawlings et al. (1994) and Meadows and Badgwell (1998).

For tracking arbitrary reference signals (rather than those generated by a finite-dimensional exogenous system $w^+ = s(w)$), variants of reference governors have been proposed. Reference governors, initially developed for control of constrained linear systems (see, for example, Gilbert, Kolmanovsky & Tan, 1994) assume that a primal controller has been designed to stabilize the plant and provide nice tracking properties *in the absence of constraints*. The reference governor then modulates the reference signal $r(t)$ to obtain a signal $g(t)$ with the following properties: $g(t)$ never causes the plant to saturate and $g(t) \rightarrow r(t)$ as $t \rightarrow \infty$. Model predictive reference governors, proposed in Bemporad, Casavola and Mosca (1997), Bemporad (1998b) and Bemporad and Mosca (1998) utilize predictions over the interval $[t, t + T]$ (t is current time) to determine $g(t)$.

5.3. Output feedback

One approach to output feedback is to use a state representation in which the state consists of current and past inputs and outputs (Mosca et al., 1990; Mosca & Zhang, 1992; Mosca, 1994; Chisci & Mosca, 1994; Bemporad et al., 1995). An alternative approach is to use an observer and certainly equivalence, i.e. replace the initial x in the optimal control problem by its estimate \hat{x} . The estimation error is an uncertainty, so robustness issues, not fully resolved, are involved though some perturbation results (Sokaert, Rawlings & Meadows, 1997) may be employed for the unconstrained case. Any state estimator may be employed (usually the Kalman filter when the system is linear). Magni, De Nicolao and Scat-

tolini (1998) establish stability of the composite system (observer and model predictive controller) when a weak detector (Vidyasagar, 1980) is employed. Michalska and Mayne (1995) establish stability of the composite system when a moving horizon observer is employed; no uncertainty exists in the model so the state can be perfectly estimated using a record of input/output data over the interval $[t - T_0, t]$ where t is the current time.

5.4. Soft constraints

While control constraints are usually hard (must be satisfied), some state constraints may be transgressed. Several researchers have studied how this extra freedom may be utilized if the state is forced into an infeasible region. In soft constraint model predictive control, violations of the state constraints are allowed, but an additional term, which penalizes constraint violations, is introduced in the objective (Ricker, Subrahmanian & Sim, 1988; Zheng & Morari, 1993, 1995; De Oliveira & Biegler, 1994; Genceli & Nikolaou, 1993; Vuthandam, Genceli & Nikolaou, 1995). Rawlings and Muske (1993) discuss an alternative approach that identifies the smallest time, which depends on the current state, beyond which the state constraint can be satisfied on an infinite horizon and enforces the state constraint only after that time. Sokaert and Rawlings (1999) analyse these approaches and propose that constraint violation at times of infeasibility be regarded as a multi-objective problem.

5.5. Adaptation

Although adaptation was one of the earliest motivations for model predictive control, a stabilizing adaptive model predictive controller for constrained systems has not yet been developed. A prior requirement for progress in this area, as in output feedback model predictive control with state estimation, is a good methodology for achieving robustness.

5.6. Optimization algorithms

Mayne (1995) provides a general discussion of the use of optimization for model predictive control. When the system is linear, the cost quadratic, and the sets \mathbb{U} and \mathbb{X} polyhedral, the optimal control problem $\mathcal{P}_N(x)$ is a quadratic programme, provided X_f is replaced, if necessary, by a polytopic inner approximation. If the system is nonlinear, $\mathcal{P}_N(x)$ is generally non-convex, so that only local, rather than global, solutions will generally be available. Then, suboptimal model predictive control may be employed although in some cases it may be possible to employ global optimization. Many optimal control algorithms have been developed; see, for example, (Polak, 1997). An overview of algorithms suitable for

solving $\mathcal{P}_N(x)$ is given in Biegler (1998). Wright (1997) and Rao, Wright and Rawlings (1998) describe algorithms developed for use in linear model predictive control.

6. Achievements and future objectives

We give below our assessment of research on model predictive control within the focus of this review, and endeavour to formulate a few desirable objectives for the near future.

6.1. Stability

Research on stability of model predictive controlled systems has now reached a relatively mature stage. The important factors for stability have been isolated and employed to develop a range of model predictive controllers that are stabilizing and differ only in their choice of the three ingredients $F(\cdot)$, X_f and $\kappa_f(\cdot)$ that are common to most forms of model predictive control. These conditions are merely sufficient and several researchers (Nevisić & Primbs, 1997) are seeking relaxations. Nor do they provide a complete answer. While they provide a means for dealing with a few target states, problems arise when the set of target states is a continuum. Determination of the ingredients for the nonlinear case involves global optimization, which may be carried out off-line when there are but a few target states, but is undesirable for the on-line computation required if the target state changes. It would be desirable to have a terminal constraint set X_f which is a neighbourhood of the continuum of target sets (zero error manifold) and a local controller $\kappa_f(\cdot)$ that maintains states in X_f . This problem has not yet been resolved.

6.2. Robustness

Progress here has not been as dramatic. While the problem has been studied and is now well understood, the outcome of the research is *conceptual* controllers that work in principle but are too complex to employ. Further research is required to develop implementable robust model predictive controllers.

6.3. Other topics

Some progress has been achieved in all of the topics mentioned in Section 5. As indicated above, further research on the tracking problem when the set of constant target states is a continuum, or when it is necessary to track a moving target generated by an exogenous system $w^+ = s(w)$, is desirable. Progress in output feedback model predictive control depends on developments in

robustness *and* state estimation. The moving horizon state estimation approach is becoming better understood (Robertson & Lee, 1996; Robertson, Lee & Rawlings, 1996). The challenge in moving horizon estimation is to utilize prior information without prejudicing the stability of the observer (Rao & Rawlings, 1998).

There are no recent significant results on adaptive model predictive control; progress here also depends on progress on the robustness issue. There is considerable activity in the area of optimal control algorithms, some of it specifically directed to use in model predictive control. Although the possibility of achieving stability through sub-optimal model predictive control has been established, there has been little exploitation of this extra freedom to develop model predictive controllers for nonlinear systems where optimality is hard to obtain.

6.4. Hybrid systems

Model predictive control has traditionally been developed for control of finite-dimensional discrete-time systems or continuous-time systems with piecewise constant control; the optimal control problem for the latter can be formulated, at least conceptually, as a discrete-time problem. However, all processes (and model predictive control has been almost exclusively employed for process control) contain discrete components such as valves, switches, speed selectors and overrides, in addition to continuous components that are described by difference or differential equations. Consideration of hybrid systems that possess both types of components model opens up a rich area of research relevant to a range of important problems such as control and supervisory schemes in the process industries (Slupphaug & Foss, 1997; Slupphaug, 1998; Bemporad & Morari, 1998,1999; Morari, Bemporad & Mignone, 1999). Many system theoretic concepts, as well as control strategies like model predictive control, require re-examination in this setting.

Appendix A

Exponential stability results if the value function $V_N^0(\cdot)$ satisfies the following conditions: there exist positive constants a, b and c such that

$$a|x|^2 \leq V_N^0(x) \leq b|x|^2, \quad \forall x \in X_N, \quad (\text{A.1})$$

$$V_N^*(x, \kappa_N(x)) \leq -c|x|^2, \quad \forall x \in X_N. \quad (\text{A.2})$$

Our assumption that $\ell(x, u) \geq c|(x, u)|^2 \geq c|x|^2$ guarantees the left-hand inequality in (A.1) (with a replacing c) and (A.2). But, as observed by Jadbabaie et al. (1999),

$V_N^0(x) \leq F(x)$ for all $x \in X_f$. Often $F(x) = (\frac{1}{2})|x|_P^2$ where $P > 0$, so that $V_N^0(x) \leq b|x|^2$ for all $x \in X_f$, some $b > 0$. Since any $x \in X_N \setminus X_f$ is steered to X_f in finite time, exponential stability results. Jadbabai et al. prove their observation as follows. Let x be an arbitrary point in X_f and let $\{x^f(k; x)\}$ and $\{u^f(k; x)\}$ be the state and control sequences resulting from initial state x and controller $\kappa_f(\cdot)$. Then, by A4

$$F(x) \geq \sum_{k=0}^{N-1} \ell(x^f(k; x), u^f(k; x)) + F(x^f(N; x)).$$

But, by optimality (since $x^f(N; x) \in X_f$ by A3)

$$V_N^0(x) \leq \sum_{k=0}^{N-1} \ell(x^f(k; x), u^f(k; x)) + F(x^f(N; x)).$$

Hence $V_N^0(x) \leq F(x)$ for all $x \in X_f$, for all $N > 0$.

References


- Alamir, M., & Bornard, G. (1995). Stability of a truncated infinite constrained receding horizon scheme: The general discrete nonlinear case. *Automatica*, 31(9), 1353–1356.
- Allgöwer, F., & Zheng, A. (1999). *Model predictive control: Assessment and future directions*. Proceedings of international workshop on model predictive control, Ascona, 1998, Berlin: Springer.
- Allwright, J. C., (1993). On min-max model-based predictive control. *Proceedings Oxford symposium on advances in model based predictive control*, Oxford (pp. 4153–426).
- Åström, K. J. (1970). *Introduction to stochastic control theory*, New York, USA: Academic Press.
- Badgwell, T. A. (1997). Robust model predictive control of stable linear systems. *International Journal of Control*, 68(4), 797–818.
- Bellman, R. (1957). *Dynamic programming*. Princeton: Princeton University Press.
- Bemporad, A. (1998a). A predictive controller with artificial Lyapunov function for linear systems with input/state constraints. *Automatica*, 34(10), 1255–1260.
- Bemporad, A. (1998b). Reference governor for constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 43(3), 415–419.
- Bemporad, A., Casavola, A., & Mosca, E. (1997). Nonlinear control of constrained linear systems via predictive reference management. *IEEE Transactions on Automatic Control*, 42(3), 340–349.
- Bemporad, A., Chisci, L., & Mosca, E. (1995). On the stabilizing property of SIORHC. *Automatica*, 30, 2013–2015.
- Bemporad, A., & Morari, M. (1998). Predictive control of constrained hybrid systems. *Preprints of international symposium on nonlinear model predictive control: Assessment and future directions*, Monte Verità, Ascona, Switzerland (pp. 108–127).
- Bemporad, A., & Morari, M. (1999). Control of systems integrating logic, dynamics, and constraints. *Automatica*, 35, 407–427.
- Bemporad, A., & Mosca, E. (1998). Fulfilling hard constraints in uncertain linear systems by reference managing. *Automatica*, 34(3), 451–461.
- Biegler, L. T. (1998). Advances in nonlinear programming concepts for process control. *Journal of Process Control*, 8(5–6), 301.
- Bitmead, R. R., Gevers, M., & Wertz, V. (1990). *Adaptive optimal control — The thinking man's GPC*. Englewood Cliffs, NJ: Prentice-Hall.
- Camacho, E., & Bordons, C. (1998). *Model predictive control*. Berlin: Springer.
- Chen, C. C., & Shaw, L. (1982). On receding horizon feedback control. *Automatica*, 18, 349–352.
- Chen, H., & Allgöwer, F. (1998a). Nonlinear model predictive control schemes with guaranteed stability. In: R. Berber, C. Kravaris (Eds.), *NATO ASI on nonlinear model based process control* (pp. 465–494). Dordrecht: Kluwer.
- Chen, H., & Allgöwer, F. (1998b). A quasi-infinite horizon nonlinear model predictive control scheme with guaranteed stability. *Automatica*, 14(10), 1205–1217.
- Chen, H., Scherer, C., & Allgöwer, F. (1997). A game theoretic approach to nonlinear robust receding horizon control of constrained systems. *Proceedings of the American control conference*, Albuquerque, NM.
- Chisci, L., Lombardi, A., & Mosca, E. (1996). Dual receding horizon control of constrained discrete-time systems. *European Journal of Control*, 2, 278–285.
- Chisci, L., & Mosca, E. (1994). Stabilizing I-O receding horizon control of CARMA plants. *IEEE Transactions on Automatic Control*, 39(3), 614–618.
- Chmielewski, D., & Manousiouthakis, V. (1996). On constrained infinite-time linear quadratic optimal control. *Systems & Control Letters*, 29, 121–129.
- Clarke, D. W. (1994). *Advances in model predictive control*. Oxford, UK: Oxford Science Publications.
- Clarke, D. W., Mohtadi, C., & Tuffs, P. S. (1987a). Generalized predictive control. Part 1: The basic algorithms. *Automatica*, 23(2), 137–148.
- Clarke, D. W., Mohtadi, C., & Tuffs, P. S. (1987b). Generalized predictive control. Part 2: Extensions and interpretations. *Automatica*, 23(2), 149–160.
- Clarke, D. W., & Scattolini, R. (1991). Constrained receding horizon predictive control. *Proceedings of the IEE, Part D, Control theory and applications*, 138, 347–354.
- Cutler, C. R., & Ramaker, B. L. (1980). Dynamic matrix control — a computer control algorithm. *Proceedings Joint Automatic Control Conference*, San Francisco, CA.
- De Keyser, R. M., & Van Cauwenberghe, A. R. (1979). A self-tuning predictor as operator guide. *Proceedings of the fifth IFAC symposium on identification and system parameter estimation* (pp. 1249–1256). Oxford: Pergamon.
- De Nicolao, G., Magnani, L., Magni, L., & Scattolini, R. (1999a). On stabilizing receding horizon control for nonlinear discrete time systems. *Proceedings of 38th IEEE conference on decision and control*.
- De Nicolao, G., Magni, L., & Scattolini, R. (1996a). On the robustness of receding horizon control with terminal constraints. *IEEE Transactions on Automatic Control*, 41, 451–453.
- De Nicolao, G., Magni, L., & Scattolini, R. (1996b). Robust predictive control of systems with uncertain impulse response. *Automatica*, 32(10), 1475–1479.
- De Nicolao, G., Magni, L., & Scattolini, R. (1996c). Stabilizing nonlinear receding horizon control via a nonquadratic penalty. *Proceedings of the IMACS multiconference CESA*, vol. 1, Lille, France (pp. 185–187).
- De Nicolao, G., Magni, L., & Scattolini, R. (1998). Stability and robustness of nonlinear receding horizon control. *International symposium on nonlinear model predictive control: Assessment and future directions*, Ascona.
- De Nicolao, G., Magni, L., & Scattolini, R. (1999b). Robustness of receding horizon control for nonlinear discrete-time systems. *Robustness in identification and control*. Berlin: Springer, to appear.
- De Nicolao, G., & Sepulchre, R. (1997). Fake Riccati equations for stable receding-horizon control. *Proceedings of ECC conference*.
- De Oliveira, N. M. C., & Biegler, L. T. (1994). Constraint handling and stability properties of model-predictive control. *A.I.Ch.E. Journal*, 40(7), 1138–1155.

- De Oliveira, S. L., Nevistic, V., & Morari, M. (1995). Control of nonlinear systems subject to input constraints. *IFAC symposium on nonlinear control system design, Tahoe City, CA* (pp. 15–20).
- Fleming, W. H., & Rishel, R. W. (1975). *Deterministic and stochastic optimal control*. New York: Springer.
- Gallestey, E., & James, M. R. (1999). H_∞ control of nonlinear systems using receding horizon controllers, in preparation.
- García, C. E., & Morshedi, A. M. (1986). Quadratic programming solution of dynamic matrix control (QDMC). *Chemical Engineering Communications*, 46, 73–87.
- García, C. E., Prett, D. M., & Morari, M. (1989). Model predictive control: theory and practice — a survey. *Automatica*, 25(3), 335–348.
- Gauthier, J. P., & Bornard, G. (1983). Commande multivariable en présence de contraintes de type inégalité. *Revue d'Automatique d'Informatique et de Recherche Opérationnelle (RAIRO)*, 17(3), 205–222.
- Genceli, H., & Nikolaou, M. (1993). Robust stability analysis of constrained l_1 norm model predictive control. *A.I.Ch.E. Journal*, 39(12), 1954–1965.
- Gilbert, E. C., Kolmanovsky, I., & Tan, K. T. (1994). Nonlinear control of discrete-time linear systems with state and control constraints: A reference governor with global convergence properties. *Proceedings of the 33rd IEEE conference on decision and control*. Lake Buena Vista, FL (pp. 144–149).
- Gilbert, E. G., & Tan, K. T. (1991). Linear systems with state and control constraints: the theory and application of maximal output admissible sets. *IEEE Transactions on Automatic Control*, AC-36, 1008–1020.
- Jadbabaie, A., Yu, J., & Hauser, J. (1999). Unconstrained receding horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, submitted.
- Kalman, R. E. (1960). Contributions to the theory of optimal control. *Boletín Sociedad Matemática Mexicana*, 5, 102–119.
- Kantor, J. C., García, C. E., & Carnahan, B. (1997). (Eds.), *Fifth international conference on chemical process control, CACHE, A.I.Ch.E.*
- Keerthi, S. S. (1986). *Optimal feedback control of discrete-time systems with state-control constraints and general cost functions*. Ph.D. thesis, University of Michigan.
- 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 Application*, 57, 265–293.
- Kleinman, B. L. (1970). An easy way to stabilize a linear constant system. *IEEE Transactions on Automatic Control*, 15(12), 693.
- Kothare, M. V., Balakrishnan, V., & Morari, M. (1996). Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10), 1361–1379.
- Kurtz, M. J., & Henson, M. A. (1997). Input–output linearizing control of constrained nonlinear processes. *Journal of Process Control*, 7(1), 3–17.
- Kwon, W. H., Bruckstein, A. M., & Kailath, T. (1983). Stabilizing state-feedback design via the moving horizon method. *International Journal of Control*, 37(3), 631–643.
- Kwon, W. H., & Pearson, A. E. (1977). A modified quadratic cost problem and feedback stabilization of a linear system. *IEEE Transactions on Automatic Control*, 22(5), 838–842.
- Lall, S., & Glover, K. (1994). A game theoretic approach to moving horizon control. In: D. W. Clarke, *Advances in model-based predictive control* (pp. 131–144). Oxford: Oxford University Press.
- Lee, E. B., & Markus, L. (1967). *Foundations of optimal control theory*. New York: Wiley.
- Lee, J. H., & Cooley, B. (1997). Recent advances in model predictive control and other related areas. In: J. C. Kantor, C.E. García, & B. Carnahan (Eds.), *Fifth international conference on chemical process control, CACHE, AICHE* (pp. 201–216).
- Lee, J. H., Morari, M., & García, C. E. (1994). State-space interpretation of model predictive control. *Automatica*, 30(4), 707–717.
- Lee, J. H., & 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. (1998). Output feedback receding-horizon control of discrete-time nonlinear systems. *Proceedings of the IFAC nonlinear control systems design symposium, Enschede, The Netherlands*.
- Magni, M., De Nicolao, G., Scattolini, R., & Allgöwer, F. (1999a). H_∞ receding horizon control for non-linear discrete-time systems. *Systems & Control Letters*, submitted.
- Magni, L., Nijmeijer, H., & van der Schaft, A. (1999b). A receding horizon approach to the nonlinear H_∞ problem. *Automatica*, submitted.
- Magni, L., & Sepulchre, R. (1997). Stability margins of nonlinear receding-horizon control via inverse optimality. *Systems & Control Letters*, 32, 241–245.
- Marquis, P., & Broustail, J. P. (1988). SMOC, a bridge between state space and model predictive controllers: application to the automation of a hydrotreating unit. In: T. J. McAvoy, Y. Arkun, & E. Zafriou (Eds.), *Proceedings of the IFAC workshop on model based process control* (pp. 37–43). Oxford: Pergamon Press.
- Mayne, D. Q. (1995). Optimization in model based control, *Proceedings of the IFAC symposium on dynamics and control chemical reactors and batch processes (Dycord + '95), Helsingor, Denmark* (pp. 229–242). Oxford: Elsevier Science. Plenary address.
- Mayne, D. Q. (1997). Nonlinear model predictive control: An assessment. In: J. C. Kantor, C. E. García, & B. Carnahan (Eds.), *Fifth International Conference on Chemical Process Control, CACHE, A.I.Ch.E.* (pp. 217–231).
- Mayne, D. Q., & Michalska, H. (1990). Receding horizon control of non-linear systems. *IEEE Transactions on Automatic Control*, 35(5), 814–824.
- Meadows, E. S. (1997). Dynamic programming and model predictive control. *Proceedings of American Control Conference, Albuquerque* (pp. 1635–1639).
- Meadows, E. S., & Badgwell, T. A. (1998). Feedback through steady-state target optimization for nonlinear model predictive control. *Journal of Vibration and Control*, 4, 61–74.
- Meadows, E. S., Henson, M. A., Eaton, J. W., & Rawlings, J. B. (1995). Receding horizon control and discontinuous state feedback stabilization. *International Journal of Control*, 62, 1217–1229.
- Michalska, H. (1997). A new formulation of receding horizon control without a terminal constraint on the state. *European Journal of Control*, 3(1), 2–14.
- Michalska, H., & Mayne, D. Q. (1993). Robust receding horizon control of constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 38, 1623–1632.
- Michalska, H., & Mayne, D. Q. (1995). Moving horizon observers and observer-based control. *IEEE Transactions on Automatic Control*, 40(6), 995–1006.
- Morari, M., Bemporad, A., & Mignone, D. (1999). A framework for control, state estimation and verification of hybrid systems. *Automatisierungstechnik*, to appear.
- Morari, M., & De Oliveira, S. L. (1998). Contractive model predictive control for constrained nonlinear systems. *IEEE Transactions on Automatic Control*, in press.
- Morari, M., & Lee, J. H. (1999). Model predictive control: Past, present and future. *Computers and Chemical Engineering*, 23, 667–682.
- Mosca, E. (1994). *Optimal, predictive and adaptive control*, Information and system science Series. Englewoods Cliffs, NJ: Prentice Hall.
- Mosca, E., Lemos, J. M., & Zhang, J. (1990). Stabilizing I/O receding horizon control. *Proceedings 29th IEEE conference on decision and control, Honolulu* (pp. 2518–2523).
- Mosca, E., & Zhang, J. (1992). Stable redesign of predictive control. *Automatica*, 28(6), 1229–1233.

- Muske, K. R., & Rawlings, J. B. (1993). Model predictive control with linear models. *A.I.Ch.E. Journal*, 39(2), 262–287.
- Nevistić, V., & Primbs, J. A. (1997). *Finite receding horizon linear quadratic control: A unifying theory for stability and performance analysis*. Technical Report CIT-CDS 97-001, California Institute of Technology, Pasadena, CA.
- Parisini, T., & Zoppoli, R. (1995). A receding horizon regulator for nonlinear systems and a neural approximation. *Automatica*, 31(10), 1443–1451.
- Peterka, V. (1984). Predictor-based self tuning control. *Automatica*, 20, 39–50.
- Polak, E. (1997). *Optimization: Algorithms and consistent approximations*. New York: Springer, ISBN 0-387-94971-2.
- Polak, E., & Yang, T. H. (1993a). Moving horizon control of linear systems with input saturation and plant uncertainty — Part 1: Robustness. *International Journal of Control*, 58(3), 613–638.
- Polak, E., & Yang, T. H. (1993b). Moving horizon control of linear systems with input saturation and plant uncertainty-Part 2: Disturbance rejection and tracking. *International Journal of Control*, 58(3), 639–663.
- Poubelle, M. A., Bitmead, R. R., & Gevers, M. (1988). Fake algebraic Riccati techniques and stability. *IEEE Transactions on Automatic Control*, AC-31, 379–381.
- Prett, D. M., & Gillette, R. D. (1980). Optimization and constrained multivariable control of a catalytic cracking unit. *Proceedings of the joint automatic control conference*, WP5-c.
- Primbs, J. A., & Nevistić, V. (1997). *Constrained finite receding horizon linear quadratic control*. Technical Report CIT-CDS 97-002, California Institute of Technology, Pasadena, CA.
- Primbs, J. A., Nevistić, V., & Doyle, J. C. (1998). On receding horizon extensions and control Lyapunov functions. *Proceedings of the American automatic control conference*, pp. 3276–3280.
- Propoi, A. I. (1963). Use of linear programming methods for synthesizing sampled-data automatic systems. *Automation and Remote Control*, 24(7), 837–844.
- Qin, S. J., & Badgwell, T. A. (1997). An overview of industrial model predictive control technology. In: J. C. Kantor, C. E. Garcia, & B. Carnahan (Eds.), *Fifth International Conference on Chemical Process Control*, CACHE, AIChE, (pp. 232–256).
- Rao, C. V., & Rawlings, J. B. (1998). Nonlinear moving horizon estimation. *Preprints of international symposium on nonlinear model predictive control: Assessment and future directions*, Monte Verità, Ascona, Switzerland (pp. 146–163).
- Rao, C. V., Wright, S. J., & Rawlings, J. B. (1998). Application of interior-point methods to model predictive control. *Journal of Optimization Theory and Applications*, 99(3), 723–757.
- Rawlings, J. B., Meadows, E. S., & Muske, K. R. (1994). Nonlinear model predictive control: A tutorial and survey. *ADCHEM'94 Proceedings, Kyoto, Japan* (pp. 185–197).
- Rawlings, J. B., & Muske, K. R. (1993). Stability of constrained receding horizon control. *IEEE Transactions on Automatic Control*, AC-38(10), 1512–1516.
- Richalet, J., Rault, A., Testud, J. L., & Papon, J. (1976). Algorithmic control of industrial processes. *Proceedings of the Fourth IFAC symposium on identification and system parameter estimation* (pp. 1119–1167).
- Richalet, J., Rault, A., Testud, J. L., & Papon, J. (1978). Model predictive heuristic control: Applications to industrial processes. *Automatica*, 14, 413–428.
- Ricker, N. L., Subrahmanian, T., & Sim, T. (1988). Case studies of model predictive control in pulp and paper production. *Proceedings of the 1988 IFAC workshop on model based predictive control*. Oxford: Pergamon Press (pp. 13–22).
- Robertson, D. G., Lee, J. H., & Rawlings, J. B. (1996). A moving horizon-based approach for least squares state estimation. *A.I.Ch.E. Journal*, 42(8), 2209–2224.
- Robertson, D., & Lee, J. (1996). Statistical interpretation of constrained moving horizon estimation. *Proceedings of A.I.Ch.E. annual meeting*, Chicago, IL.
- Rossiter, J. A., Kouvaritakis, B., & Rice, M. J. (1998). A numerically robust state-space approach to stable-predictive control strategies. *Automatica*, 34(1), 65–74.
- Scokaert, P. O. M., & Mayne, D. Q. (1998). Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8), 1136–1142.
- Scokaert, P. O. M., Mayne, D. Q., & Rawlings, J. B. (1999). Suboptimal model predictive control (feasibility implies stability). *IEEE Transactions on Automatic Control*, 44(3), 648–654.
- Scokaert, P. O. M., & Rawlings, J. B. (1998). Constrained linear quadratic regulation. *IEEE Transactions on Automatic Control*, 43(8), 1163–1169.
- Scokaert, P. O. M., Rawlings, J. B., & Meadows, E. S. (1997). Discrete-time stability with perturbations: Applications to model predictive control. *Automatica*, 33(3), 463–470.
- Scokaert, P. O., & Rawlings, J. B. (1999). Feasibility issues in linear model predictive control. *A.I.Ch.E. Journal*, 45(8), 1649–1659.
- Slupphaug, O. (1998). *On robust constrained nonlinear control and hybrid control BMI and MPC based state-feedback schemes*, Ph.D. Norwegian University of Science and Technology.
- Slupphaug, O., & Foss, B. A. (1997). Model predictive control for a class of hybrid systems. *Proceedings of European control conference*, Brussels, Belgium.
- Soeterboek, R. (1992). *Predictive control — a unified approach*. Englewood Cliffs, NJ: Prentice-Hall.
- Sznaier, M., & Damborg, M. J. (1987). Suboptimal control of linear systems with state and control inequality constraints. *Proceedings of the 26th IEEE conference on decision and control*, Los Angeles (pp. 761–762).
- Sznaier, M., & Damborg, M. J. (1990). Heuristically enhanced feedback control of constrained discrete-time linear systems. *Automatica*, 26, 521–532.
- Tadmor, G. (1992). Receding horizon revisited: An easy way to stabilize an LTV system. *Systems & Control Letters*, 8, 285–294.
- Thomas, Y. A. (1975). Linear quadratic optimal estimation and control with receding horizon. *Electronics Letters*, 11, 19–21.
- van der Schaft, A. J. (1996). *L₂-gain and passivity techniques in nonlinear control*. Berlin: Springer.
- Vidyasagar, M. (1980). On the stabilization of nonlinear systems using state detection. *IEEE Transactions on Automatic Control*, 25, 504–509.
- Vuthandam, P., Genceli, H., & Nikolaou, M. (1995). Performance bounds for robust quadratic dynamic matrix control with end condition. *A.I.Ch.E. Journal*, 41(9), 2083–2097.
- Wright, S. J. (1997). Applying new optimization algorithms to model predictive control. In: J.C. Kantor, C.E. Garcia, B. Carnahan (Eds.), *Fifth international conference on chemical process control*, CACHE, AIChE (pp. 147–155).
- Yang, T. H., & Polak, E. (1993). Moving horizon control of nonlinear systems with input saturation, disturbances and plant uncertainty. *International Journal of Control*, 58, 875–903.
- Ydstie, B. E. (1984). Extended horizon adaptive control. *Proceedings of the Ninth world congress of IFAC*, Budapest, Hungary.
- Zheng, A. (1997). A computationally efficient nonlinear MPC algorithm. *Proceedings of the American control conference*, Albuquerque, NM (pp. 1623–1627).
- Zheng, A., & Morari, M. (1993). Robust stability of constrained model predictive control. *Proceedings of the American control conference* (pp. 379–383).
- Zheng, A., & Morari, M. (1995). Stability of model predictive control with mixed constraints. *IEEE Transactions on Automatic Control*, 40(10), 1818–1823.



David Mayne was born in South Africa and has held appointments at the University of the Witwatersrand, South Africa, the University of London, England, and the University of California, Davis. He received the degrees of Ph.D. and D.Sc. from the University of London and the degree of Doctor of Technology, *honoris causa*, from the University of Lund, Sweden. His research interests include optimization, model predictive control, nonlinear control and adaptive control. He is a Fellow of the Royal Society and a Fellow of the Royal Academy of Engineering.



James B. Rawlings was born in Gary, Indiana, USA in 1957. He received the B.S. from the University of Texas in 1979 and the Ph.D. from the University of Wisconsin in 1985, both in Chemical Engineering. He spent one year at the University of Stuttgart as a NATO postdoctoral fellow and then joined the faculty at the University of Texas. He moved to the University of Wisconsin in 1995 and is currently the Paul A. Elfers Professor of Chemical Engineering and the co-director of the

Texas–Wisconsin Modeling and Control Consortium (TWMCC). His research interests are in the areas of chemical process modeling, monitoring and control, nonlinear model predictive control, moving horizon state estimation, particulate systems modeling, and crystallization.



Christopher V. Rao received the B.S. degree in Chemical Engineering from Carnegie Mellon University, Pittsburgh, PA, in May 1994. He is now a Ph.D. candidate in Chemical Engineering at the University of Wisconsin–Madison. His research interests include optimization, process control, computational molecular biology, and metabolic engineering.



Pierre Scokaert received the B.Eng. degree from the University of Reading in 1990, and the D.Phil. degree from the University of Oxford in 1994. During the period of 1994 to 1995, he held the post of post-doctoral researcher in Department of Chemical Engineering at the University of Texas at Austin, and at the University of Wisconsin at Madison. In 1996, he then held the post of post-doctoral researcher in the Department of Computer Sciences and Electrical Engineering at the University of

California at Davis. He is currently employed in France, as a Research Engineer at the Centre National d'Etudes des Telecommunications, France Telecom's research center.