

Hybrid Predictive Control of Process Systems

Prashant Mhaskar, Nael H. El-Farra, and Panagiotis D. Christofides

Dept. of Chemical Engineering, University of California, Los Angeles, CA 90095

DOI 10.1002/aic.10115

Published online in Wiley InterScience (www.interscience.wiley.com).

The operation of chemical processes often requires respecting constraints on manipulated inputs and process states. Input constraints typically reflect limits on the capacity of control actuators, such as valves or pumps, whereas state constraints represent desirable ranges of operation for process variables, such as temperatures or concentrations. Constraints, however, limit the set of initial conditions, starting from where a process can be stabilized at a possibly open-loop unstable steady state. Therefore, in control of constrained processes, it is important to obtain an explicit characterization of the region of closed-loop stability. Model predictive control (MPC) provides a suitable framework for implementing control that respects manipulated input and process variable constraints while meeting prescribed performance objectives. Unfortunately, the implicit nature of the feedback law in MPC (the control action is computed by solving on-line a constrained optimization problem at each sampling time) makes the a priori computation of the closed-loop stability region a very difficult task. Such a computation, however, is possible when Lyapunov-based bounded control techniques are used to design controllers for the stabilization of systems with manipulated input constraints. Motivated by this, the present work proposes a hybrid predictive control structure, which seamlessly unites MPC and bounded control, for output feedback stabilization of linear systems with input constraints. The design of the proposed structure is based on the idea of using a bounded controller, with its associated region of stability, as a “fallback” controller for the output feedback implementation of MPC. Switching laws are derived to orchestrate the transition between the two controllers in a way that reconciles their respective stability and optimality properties. The switched closed-loop system is shown to possess a region of guaranteed stability, achieved by appropriately switching between the two controllers. The proposed hybrid control structure is shown to provide a safety net for the practical implementation of MPC by providing a priori knowledge, through off-line computations, of a large set of initial conditions for which closed-loop stability is guaranteed. Finally, two simulation studies are presented to demonstrate the implementation and evaluate the effectiveness of the proposed hybrid predictive control scheme, as well as test its robustness with respect to modeling errors and measurement noise. © 2004 American Institute of Chemical Engineers AIChE J, 50: 1242–1259, 2004

Keywords: *input constraints, model predictive control, Lyapunov-based bounded control, state estimation, controller switching, closed-loop stability region, process control*

Introduction

The operation of chemical processes invariably requires respecting constraints. The physical limitations on the func-

tioning of the control actuators (for example, the maximum flow rate that a given pump can generate is always limited) impose constraints on the manipulated inputs, whereas performance, safety, and environmental considerations (requiring, for example, the reactor temperature to stay below a certain value, or the emissions to be within an allowable range) give rise to state constraints. Constraints automatically impose limitations on our ability to steer the dynamics of the closed-loop system

Correspondence concerning this article should be addressed to P. D. Christofides at pd@seas.ucla.edu.

at will, and can cause severe deterioration in the closed-loop performance, and may even lead to closed-loop instability, if not explicitly taken into account at the stage of controller design.

For open-loop unstable processes, one of the key limitations imposed by input constraints is the restriction on the set of initial states of the closed-loop system that can be steered to the origin with the available control action (the so-called null controllable region). Furthermore, a given control law can typically guarantee closed-loop stability from only a subset of initial states included in the null controllable region. An adequate characterization of the stability properties of a given stabilizing control law involves a priori (that is, before controller implementation) knowledge of this region of stability, required to ascertain stabilizability of the system from a given initial condition. The absence of an a priori explicit characterization of this set (or an appropriate estimate thereof) can have an impact on the practical implementation of the given control policy.

Model predictive control (MPC), also known as receding horizon control, is one control method for handling constraints (both on manipulated inputs and state variables) within an optimal control setting (Richalet et al., 1978). In MPC, the control action is obtained by solving repeatedly, on-line, a finite-horizon constrained open-loop optimal control problem. The popularity of this approach stems largely from its ability to handle, among other issues, multivariable interactions, constraints on controls and states, and optimization requirements. Its success in many commercial applications is well documented in the literature (see, for example, Garcia et al., 1989; Qin and Badgwell, 1997). Numerous research studies have investigated the stability properties of model predictive controllers and led to a plethora of MPC formulations that focus on closed-loop stability (see, for example, De Oliveira and Biegler, 1994; Genceli and Nikolaou, 1993; Keerthi and Gilbert, 1988; Rawlings and Muske, 1993; see also the review paper by Mayne et al., 2000).

The significant progress in understanding and improving the stability properties of MPC notwithstanding, the issue of obtaining, a priori, an explicit characterization of the region of constrained closed-loop stability for model predictive controllers remains to be adequately addressed. Part of the difficulty in this direction is attributed to the fact that the stability of model predictive controllers depends on a complex interplay between several factors such as the choice of the horizon length, the penalties in the performance index, the initial condition, and the constraints on the state variables and manipulated inputs. A priori knowledge of the stability region requires an explicit characterization of these interplays, which is a very difficult task. This difficulty can have an impact on the practical implementation of MPC by imposing the need for extensive closed-loop simulations over the whole set of possible initial conditions to check for closed-loop stability, or by potentially limiting operation within an unnecessarily small neighborhood of the nominal equilibrium point.

The desire to implement control approaches that allow for an explicit characterization of their stability properties has motivated significant work on the design of stabilizing control laws, using Lyapunov techniques, that provide explicitly defined, large regions of attraction for the closed-loop system; the reader may refer to Kokotovic and Arcak (2001) for a survey of




Figure 1. Hybrid control structure.

results in this area. Recently, in El-Farra and Christofides (2001a,b) and El-Farra and Christofides (2003a), a class of Lyapunov-based bounded robust nonlinear controllers, inspired by the results on bounded control originally presented in Lin and Sontag (1991), was developed. The controllers enforce robust stability in the closed-loop system and provide, at the same time, an explicit characterization of the region of guaranteed closed-loop stability. Despite their well-characterized stability and constraint-handling properties, those Lyapunov-based controllers are not guaranteed to be optimal with respect to an arbitrary performance criterion.

However, the fact that Lyapunov-based control methods provide an easily implementable controller with an explicitly characterized stability region for the closed-loop system motivates developing control techniques that use the a priori guarantees of stability provided by the bounded controller as a safety net for the implementation of the high performance model predictive controller. Reconciliation of different controllers (designed to satisfy different control objectives) calls for invoking the hybrid control paradigm, where the control structure consists of a blend of continuous (that is, classical) controllers and discrete components (for example, a logic-based supervisor that determines switching between the various continuous controllers; see Figure 1). In El-Farra et al. (2004b), we proposed such a hybrid control structure that uses switching between a model predictive controller and a bounded analytical controller for the state feedback stabilization of linear systems with input constraints. The hybrid control structure was subsequently extended to nonlinear systems with input constraints in El-Farra et al. (2004a). This structure differs from other hybrid control structures found in the literature, in the sense that it uses *structurally different* controllers as opposed to most of the work on hybrid control that uses switching between different models (which results in an implicit switching between different controllers), or uses multiple *structurally similar* controllers; examples along these lines include gain scheduling as a tool for control of nonlinear systems (for example, see Rugh and Shamma, 2000), the use of multiple linear models for transition control (for example, see Banerjee and Arkun, 1998; Sun and Hoo, 1999) and for predictive control (for example, see Aufderheide et al., 2001; Schott and Bequette, 1998), and hybrid control of switched nonlinear systems (El-Farra and Christofides, 2003b).

In addition to constraints, another important issue in chemical process control is the lack of complete measurements of the process state variables. This problem is usually addressed

by designing appropriate state estimators (observers) to obtain estimates of the state variables from the available measurements (outputs). The state estimators are then appropriately coupled with state feedback controllers, leading to output feedback controllers that can be implemented with the available measurements. The problems of state estimation and output feedback controller design for constrained systems have been the subject of several research studies. For example, recent results on moving horizon estimation, as an extension of Kalman filtering, for constrained and nonlinear processes, can be found in Rao and Rawlings (2002). Examples of results on output feedback stabilization of constrained systems include output feedback control of linear systems (Shamma and Tu, 1998) and the combination of a moving horizon regulator with a moving horizon observer for output feedback control of nonlinear systems (Michalska and Mayne, 1995). In these works, however, the output feedback stability region of the constrained closed-loop system is not explicitly characterized.

Motivated by the above considerations, the present work proposes a hybrid predictive control structure that seamlessly unites MPC and bounded control for the output feedback stabilization of linear systems with input constraints. The proposed strategy is based on the idea of embedding the implementation of MPC within the output feedback stability region of the bounded controller, and using this controller as a “fall-back” in the event that MPC is unable to achieve closed-loop stability (due, for example, to improper tuning of MPC parameters). Because of the absence of complete state measurements, a state observer is constructed to provide the controllers, as well as the supervisor, with appropriate state estimates. The observer is tuned in a way so as to guarantee closed-loop stability for all initial conditions within the bounded controller’s output feedback stability region (which can be chosen arbitrarily close in size to its state feedback counterpart, provided that the observer gain is sufficiently large). Switching laws, which monitor the evolution of the state estimates, are derived to orchestrate the transition between the two controllers in a way that reconciles their respective stability and optimality properties and safeguards closed-loop stability in the event of MPC infeasibility or instability. In addition to the set of switching rules being different from the one proposed in El-Farra et al. (2004b) under state feedback, an important characteristic of the hybrid predictive control strategy under output feedback is the inherent coupling, brought about by the lack of full state measurements, between the tasks of controller design, characterization of the stability region, and supervisory switching logic design, on one hand, and the task of observer design, on the other. The rest of the article is organized as follows: initially, we present some preliminaries that describe the class of systems considered, and review the design of a Luenberger state observer. Next, we review the design procedure, and stability properties, of both MPC and Lyapunov-based bounded control, respectively. Then, we formulate the hybrid control problem within the framework of switched linear systems and propose the hybrid predictive control structure. The theoretical underpinnings and practical implications of the proposed hybrid predictive control structure are highlighted, and possible extensions of the supervisory switching logic, which address a variety of practical implementation issues, are discussed. Finally, two simulation studies are presented to demonstrate the implementation and evaluate the effectiveness of

the proposed control strategy, as well as test its robustness with respect to modeling errors and measurement noise.

Preliminaries

In this work, we consider the problem of output feedback stabilization of continuous-time linear time-invariant (LTI) systems with input constraints, with the following state-space description

$$\dot{x} = Ax + Bu \quad (1)$$

$$y = Cx \quad (2)$$

$$u \in \mathcal{U} \subset \mathbb{R}^m \quad (3)$$

where $x = [x_1 \cdots x_n]^\top \in \mathbb{R}^n$ denotes the vector of state variables, $y = [y_1 \cdots y_k]^\top \in \mathbb{R}^k$ denotes the vector of output variables, and $u = [u_1 \cdots u_m]^\top \in \mathbb{R}^m$ is the vector of manipulated inputs, taking values in a compact and convex subset, $\mathcal{U} := \{u \in \mathbb{R}^m : \|u\| \leq u_{\max}\}$, which contains the origin in its interior, where $\|\cdot\|$ denotes the standard Euclidean norm of a vector and $u_{\max} > 0$ is the magnitude of actuator constraints. The matrices A , B , and C are constant $n \times n$, $n \times m$, and $k \times n$ matrices, respectively. The pair (A, B) is assumed to be controllable and the pair (C, A) is assumed to be observable. Throughout the article, the notation $\|\cdot\|_Q$ refers to the weighted norm, defined by $\|x\|_Q^2 = x'Qx$ for all $x \in \mathbb{R}^n$, where Q is a positive-definite symmetric matrix and x' denotes the transpose of x .

We will now briefly review the design of a Luenberger state observer for the system of Eqs. 1 and 2, which will be used later in the development of the hybrid predictive control structure (see Remark 1 below for a discussion on the use of other possible estimation schemes). Specifically, we consider a standard Luenberger state observer described by

$$\dot{\hat{x}} = A\hat{x} + Bu + L(y - C\hat{x}) \quad (4)$$

where $\hat{x} = [\hat{x}_1 \cdots \hat{x}_n]^\top \in \mathbb{R}^n$ denotes the vector of estimates of the state variables and L is a constant $n \times k$ matrix that multiplies the discrepancy between the actual and estimated outputs. Under the state observer of Eq. 4, the estimation error in the closed-loop system, defined as $e = x - \hat{x}$, evolves, independently of the controller, according to the following equation

$$\dot{e} = (A - LC)e \quad (5)$$

The pair (C, A) is assumed to be observable in the sense that the observer gain matrix, L , can be chosen such that the norm of the estimation error in Eq. 5 evolves according to $\|e(t)\| \leq \kappa(\beta)\|e(0)\|\exp(-\beta t)$, where $-\beta < 0$ is the largest eigenvalue of $A - LC$ and $\kappa(\beta)$ is a polynomial function of β .

Remark 1. Referring to the state observer of Eq. 4, it should be noted that the results presented in this work are not restricted to this particular class of observers. Any other observer, which allows us to control the rate of decay of the estimation error at will, can be used. Our choice of using this particular observer design is motivated by the fact that it

provides a transparent relationship between the temporal evolution of the estimation error bound and the observer parameters. For example, this design guarantees convergence of the state estimates in a way such that for larger values of β , the error decreases faster. As we discuss later (see the section on hybrid predictive control), the ability to ensure a sufficiently fast decay of the estimation error is necessary to guarantee closed-loop stability under output feedback control. This requirement or constraint on the error dynamics is present even when other estimation schemes, such as moving horizon observers, are used (for example, see Michalska and Mayne, 1995; Rao and Rawlings, 2002) to ensure closed-loop stability. For such observers, however, it is difficult in general to obtain a transparent relationship between the tunable observer parameters and the error decay rate.

Remark 2. For large values of β , the estimation error could possibly increase to large values before eventually decaying to zero (a phenomenon known as “peaking”; see Sussmann and Kokotovic, 1991). However, this does not pose a problem in our design because the physical constraints on the manipulated inputs prevent transmission of the incorrect estimates to the process (see also Remark 9 for a detailed discussion of this issue). In the special case when C is a square matrix of rank n , the observer gain matrix L can be chosen as $L = AC^{-1} - RC^{-1}$, where R is a diagonal matrix whose diagonal elements are of the form $R_{ii} = -\beta a_i$, where $a_i \neq 0, a_i \geq 1$. For this choice of L , the evolution of each error state is completely decoupled from the rest of the error states, that is, each error state evolves according to $\dot{e}_i = -\beta a_i e_i$. In this case, no peaking occurs because each error term decreases monotonically.

Model Predictive Control

We consider model predictive control of the system described by Eq. 1, subject to the control constraints of Eq. 3. The control action at time t is conventionally obtained by solving, on-line, a finite horizon optimal control problem (Mayne, 1997) of the form

$$P(x, t) : \min\{J[x, t, u(\cdot)] | u(\cdot) \in S\} \quad (6)$$

where $S = S(t, T)$ is the family of piecewise continuous functions (functions continuous from the right), with period Δ , mapping $[t, t + T]$ into $\mathcal{U}$ and T is the specified horizon. A control $u(\cdot)$ in S is characterized by the sequence $\{u[k]\}$, where $u[k] := u(k\Delta)$. A control $u(\cdot)$ in S satisfies $u(t) = u[k]$ for all $t \in [k\Delta, (k + 1)\Delta)$. The performance index is given by

$$J[x, t, u(\cdot)] = \int_t^{t+T} [\|x''(s; x, t)\|_Q^2 + \|u(s)\|_R^2] ds + F[x(t + T)] \quad (7)$$

where R and Q are strictly positive-definite, symmetric matrices, $x''(s; x, t)$ denotes the solution of Eq. 1, due to control u , with initial state x at time t , and $F(\cdot)$ denotes the terminal penalty. In addition to penalties on the state and control action, the objective function may also include penalties on the rate of change of the input, reflecting limitations on actuator speed (for example, a large valve requiring few seconds to change posi-

tion). The minimizing control $u^0(\cdot) \in S$ is then applied to the process over the interval $[k\Delta, (k + 1)\Delta)$ and the procedure is repeated indefinitely. This defines an implicit model predictive control law

$$M(x) := u^0(t; x, t) \quad (8)$$

It is well known that the control law defined by Eqs. 6–8 is not necessarily stabilizing (Bitmead et al., 1990). To achieve closed-loop stability, early versions of MPC focused on tuning the horizon length T , and/or increasing the terminal penalty (for a survey of these approaches, see Bitmead et al., 1990), whereas more recent formulations focused on imposing stability constraints on the optimization (for surveys of different constraints proposed in the literature and the concomitant theoretical issues, see Mayne et al., 2000) or using off-line computations to generate explicit model predictive control laws (for example, see Pistikopoulos et al., 2002). The additional stability constraints serve either to enforce convergence of the states of the closed-loop system to the equilibrium point, or to force the states to reach some invariant terminal set at the end of the horizon.

For the sake of a concrete illustration of the hybrid predictive control strategy, we will consider the following stabilizing MPC formulation (Keerthi and Gilbert, 1988; Michalska and Mayne, 1990), which includes a terminal equality constraint in the optimization problem (see Remark 3 below for more advanced MPC formulations that can also be used). This model predictive controller is mathematically described by

$$J_s[x, t, u(\cdot)] = \int_t^{t+T} [x'(s)Qx(s) + u'(s)Ru(s)] ds \quad (9)$$

$$u(\cdot) = \operatorname{argmin}\{J_s[x, t, u(\cdot)] | u(\cdot) \in S\} := M_s(x)$$

$$\begin{aligned} \text{s.t. } \dot{x} &= Ax + Bu & x(0) &= x_0 \\ x(t + T) &= 0 \end{aligned} \quad (10)$$

Remark 3. It should be noted that the hybrid predictive control strategy to be proposed in this work is not restricted to the MPC formulation given in Eqs. 9 and 10, or any other MPC formulation for that matter. The results apply to any MPC formulation for which a priori knowledge of the set of admissible initial conditions is lacking (or computationally expensive to obtain). In this sense, the above formulation is really intended as a symbolic example of stabilizing MPC formulations, with the understanding that other advanced MPC formulations that use less-restrictive constraints and/or penalties on the state at the end of the horizon can be used. Examples include formulations with terminal inequality constraints, such as contractive MPC (for example, see Polak and Yang, 1993; Kothare and Morari, 2000) and CLF-based receding horizon control (for example, see Primbs et al., 2000), as well as formulations that employ a combination of terminal penalties and inequality constraints (for example, see Chen and Allgöwer, 1998; Magni et al., 2001; Sznajder et al., 2003). The reader is referred to Mayne et al. (2000) for a more exhaustive list.

Remark 4. By incorporating stability conditions directly as part of the optimization problem in Eq. 10, asymptotic stability

under state feedback MPC is guaranteed provided that the initial condition is chosen so that the optimization yields a feasible solution. However, the implicit nature of the MPC control law, obtained through repeated on-line optimization, limits our ability to obtain, a priori, an explicit characterization of the admissible initial conditions starting from where the given model predictive controller (with fixed horizon length) is guaranteed to be feasible and enforce asymptotic stability. This set is a complex function of the constraints and the horizon length. Estimates of sufficiently large horizon lengths that ensure stability (for example, see Chmielewski and Manou-siouthakis, 1996) are typically conservative and, if used, may lead to a significant computational burden due to the increased size of the optimization problem. Therefore, in practice, the initial conditions and/or horizon lengths are usually chosen using ad hoc criteria and tested through closed-loop simulations, which can add to the computational burden in implementing the model predictive controller. The difficulties encountered in characterizing the stability region under state feedback control carry over to the case of output feedback control, where the lack of state measurements requires that the control action be computed using the state estimates. Feasibility of the MPC optimization problem based on the state estimates, however, does not guarantee closed-loop stability or even the continued feasibility of the optimization problem based on the ensuing state estimates. This motivates implementing MPC within a hybrid control structure that provides a “fallback” controller for which a region of constrained closed-loop stability can be obtained off-line. Lyapunov-based controller design techniques provide a natural framework for the design of a stabilizing “fallback” controller for which an explicit characterization of the region of closed-loop stability can be obtained.

Lyapunov-Based Bounded Control

In this section, we initially assume that measurements of all the state variables are available and design a state feedback bounded controller using Lyapunov techniques and provide an explicit characterization of the stability region of the resulting closed-loop system. Then, we couple the state feedback controller with the state observer of Eq. 4 to design an output feedback controller and characterize the stability properties of the resulting closed-loop system including the set of initial conditions for which closed-loop stability under output feedback is guaranteed.

State feedback controller design

Consider the Lyapunov function candidate $V = x'Px$, where P is a positive-definite symmetric matrix that satisfies the Riccati equation:

$$A'P + PA - PBB'P = -\bar{Q} \quad (11)$$

for some positive-definite matrix $\bar{Q}$. Using this Lyapunov function, we can construct, using a modification of Sontag's formula for bounded controls proposed in Lin and Sontag (1991) (see also El-Farra and Christofides, 2003a), the following bounded nonlinear controller

$$u(x) = -2k(x)B'Px := b(x) \quad (12)$$

where

$$k(x) = \left\{ \frac{L_f^*V + \sqrt{(L_f^*V)^2 + [u_{\max}\|(L_gV)'\|^4]}}{\|(L_gV)'\|^2[1 + \sqrt{1 + [u_{\max}\|(L_gV)'\|^2]}} \right\} \quad (13)$$

with $L_f^* = x'(A'P + PA)x + \rho x'Px$, $(L_gV)' = 2B'Px$ and $\rho > 0$. This controller is continuous everywhere in the state-space and smooth away from the origin. Using a Lyapunov argument, one can show that whenever the closed-loop state trajectory evolves within the state-space region described by the set

$$\Phi(u_{\max}) = \{x \in \mathbb{R}^n : L_f^*V < u_{\max}\|(L_gV)'\|\} \quad (14)$$

The resulting control action respects the constraints (that is, $\|u\| \leq u_{\max}$) and enforces, simultaneously, the negative definiteness of the time derivative of V along the trajectories of the closed-loop system. Note that the size of the set Φ depends on the magnitude of the constraints in a way such that the tighter the constraints, the smaller the region described by this set. Starting from any initial state within $\Phi(u_{\max})$, asymptotic stability of the constrained closed-loop system can be guaranteed, provided that the closed-loop trajectory remains within the region described by $\Phi(u_{\max})$. To ensure this, we consider initial conditions that belong to an invariant subset, preferably the largest, which we denote by $\Omega(u_{\max})$. A similar idea was used in El-Farra and Christofides (2001) and El-Farra and Christofides (2003a) in the context of bounded robust control of constrained nonlinear systems. One way of constructing such a subset is using the level sets of V (see chapter 4 in Khalil, 1996 for details), that is

$$\Omega(u_{\max}) = \{x \in \mathbb{R}^n : x'Px \leq c_{\max}\} \quad (15)$$

where $c_{\max} > 0$ is the largest number for which all nonzero elements of $\Omega(u_{\max})$ are contained within $\Phi(u_{\max})$. The invariant region described by the set $\Omega(u_{\max})$ provides an estimate of the stability region, starting from where the origin of the constrained closed-loop system under the control law of Eqs. 12 and 13 is guaranteed to be asymptotically stable. To simplify the notation, we will suppress the dependency of the sets Ω and Φ on $u_{\max}$ in the remainder of the article.

Remark 5. The stability region under a given stabilizing control law is typically only a subset of the null controllable region (the exact computation of which remains an open research problem) and the bounded controller of Eqs. 12 and 13 is only one choice, out of many others, that provides an explicit characterization of the stability region. Furthermore, although the level sets of V provide only an estimate of the stability region under the bounded controller, less conservative estimates that capture larger portions of the null controllable region can be obtained using, for example, a combination of several Lyapunov functions (for example, see Hu and Lin, 2002; El-Farra et al., 2004a). In general, well-tuned bounded control laws provide larger estimates of the stability region than controllers designed without taking the constraints into account (for further details on this issue, see El-Farra and Christofides, 2001).

Remark 6. Referring to the optimality properties of the bounded controller, it should be noted that, within a well-defined subset of Ω , this class of controllers can be shown to be inverse optimal with respect to some meaningful, but unspecified a priori, cost functional that imposes meaningful penalties on the state and control action (for a proof and further details, see El-Farra and Christofides, 2001). However, unless this cost functional (which is determined only after controller design) coincides with the actual cost functional considered by the control system designer (this conclusion cannot be ascertained a priori), the performance of the bounded controller will not be optimal with respect to the actual cost functional. An explicit quantitative characterization of the cost incurred by implementing the bounded controller in this case is difficult to obtain. Furthermore, given an arbitrary cost functional, there is no transparent way of “tuning” the bounded controller to achieve optimality because, by design, the goal of the bounded controller is to achieve stabilization for an explicitly defined (a priori) set of initial conditions.

Remark 7. When comparing the optimality properties of the bounded and predictive controllers, it is important to realize that, in general, it is possible that the total cost of steering the system from the initial state to the origin under the bounded controller could be less than the total cost incurred under MPC. Keep in mind that MPC formulations typically consider minimizing a given cost over some finite horizon. Therefore, for a given objective function of the form of Eq. 7 for example, by solving the optimization problem once and implementing the resulting input trajectory in its entirety, the cost of taking the system from the initial state to the state at the end of the horizon under MPC is guaranteed to be smaller than the corresponding cost for the bounded controller. However, because MPC implementation involves repeated optimizations using finite horizons (that is, only the first control move of the input trajectory is implemented each time), one could in some cases end up with a higher total cost for MPC. This possibility, however, cannot be ascertained a priori (that is, before implementing both controllers and computing the total cost).

Stability properties under output feedback bounded control

The lack of state measurements necessitates the use of a state estimator to obtain estimates of the state variables. When a state estimator of the form of Eq. 4 is used, the resulting closed-loop system is composed of a cascade, between the error system and the plant, of the form

$$\begin{aligned}\dot{x} &= Ax + Bu(x - e) \\ \dot{e} &= (A - LC)e\end{aligned}\quad (16)$$

Note that the values of the states used in the controller contain errors. The state feedback stability region, therefore, is not exactly preserved under output feedback. However, by exploiting the error dynamics of Eq. 5, it is possible to recover arbitrarily large compact subsets of the state feedback stability region, provided that the poles of the observer are placed sufficiently far in the left half of the complex plane (which can be accomplished by choosing the observer gain parameter β sufficiently large). This idea is consistent with earlier results on semiglobal output feedback stabiliza-




Figure 2. Evolution of the trajectories of the closed-loop state and state estimate.

Starting from within Ω_b , the observer gain ensures that, before the state of the closed-loop system reaches the boundary of Ω , the state estimate has converged close enough to the true value of the state, so that the norm of the estimation error is below the allowable error tolerated by the controller within Ω .

tion of unconstrained systems using high-gain observers (for example, see Teel and Praly, 1994; Mahmoud and Khalil, 1996; Christofides, 2000), as well as results on semiglobal stability of nonlinear singularly perturbed systems (Christofides and Teel, 1996).

By viewing the estimation error as a dynamic perturbation to the state feedback problem, the basic idea, in constructing the output feedback controller and characterizing its stability region, is to design the observer in a way that exploits the robustness of the state feedback controller with respect to bounded estimation errors. To this end, given the state feedback stability region Ω , we initially quantify the robustness margin by deriving a bound on the estimation error, $e_m > 0$, that can be tolerated by the state feedback controller, in the sense that the closed-loop state trajectory does not escape Ω for all $\|e\| \leq e_m$. Once this bound is computed, the idea is to initialize the states and the state estimates within any subset $\Omega_b \in \Omega$ and choose a consistent observer gain matrix L (parameterized by β) to ensure that, before the states reach the boundary of Ω , the norm of the estimation error has decayed to a value less than e_m (see Figure 2 for a pictorial representation of this idea). These arguments are formalized in Propositions 1 and 2 below (the proofs are given in the Appendix).

Proposition 1. Consider the constrained LTI system of Eqs. 1–3 under the bounded control law of Eqs. 12 and 13 with $u = u(x - e)$, where e is the state measurement error. Then, there exists a positive real number, e_m , such that if $x(0) \in \Omega$ and $\|e(t)\| \leq e_m \forall t \geq 0$, then $x(t) \in \Omega \forall t \geq 0$.

Remark 8. The bounded control law of Eqs. 12 and 13 will be used in the next section to illustrate the basic idea of the proposed hybrid control scheme. Our choice of using this particular design is motivated by its explicit structure and

well-defined region of stability, which allows for an estimation of the robustness margin of the state feedback controller with respect to bounded measurement errors. However, the results of this work are not restricted to this particular choice of bounded controllers. Any other analytical bounded control law, with an explicit structure and well-defined region of robust stability with respect to bounded measurement errors, can also be used in implementing the proposed control strategy.

Proposition 2. Consider the constrained LTI system of Eqs. 1–3, the state observer of Eq. 4 and the bounded control law of Eqs. 12 and 13. Then, given any positive real number, δ_b , such that $\Omega_b = \{x \in \mathbb{R}^n : \|x\|_p^2 \leq \delta_b\} \subset \Omega$, there exists a positive real number β^* such that if $\|x(0)\|_p^2 \leq \delta_b$, $\|\hat{x}(0)\|_p^2 \leq \delta_b$ and $\beta \geq \beta^*$, the origin of the constrained closed-loop system is asymptotically stable.

Remark 9. Because the only assumption on C is that the pair (C, A) is observable, the observer states may “peak” before they converge to the true state values. This, however, does not pose a problem in our design because: (a) the physical constraints on the manipulated input eliminate the occurrence of instability attributed to peaking of the state estimates (that is, they prevent transmission of peaking to the process), and (b) by “stepping back” from the boundary of the state feedback stability region and choosing an appropriate value for β , the design ensures that the system states cannot leave the stability region of the bounded controller before the estimation errors have gone below the permissible value.

Remark 10. In principle, the stability region under output feedback, Ω_b , can be chosen as close as desired to Ω by increasing the observer gain parameter β . However, it is well known that large observer gains can amplify measurement noise and induce poor performance (see the section on simulation studies for how this issue is addressed in observer implementation). This points to a fundamental trade-off that cannot be resolved by simply changing the estimation scheme. For example, although one could replace the high-gain observer design with other observer designs (for example, a moving horizon estimator) to gain a better handle on measurement noise, it is difficult in such schemes to obtain an explicit relationship between the observer tuning parameters and the output feedback stability region.

Remark 11. Note that, for the bounded controller, initializing x and $\hat{x}$ within Ω_b guarantees that x cannot escape Ω because of the inherent robustness of the bounded controller design with respect to bounded estimation errors (see the proof of Proposition 2 in the Appendix). For MPC, however, there is no guarantee that x will stay in Ω , even if both x and $\hat{x}$ are initialized within Ω_b , and, therefore, whenever MPC is used, it will be necessary to rely on the evolution of the estimate, $\hat{x}(t)$, to deduce whether $x(t)$ is or is not within Ω . As we discuss in greater detail in the next section, the ability to reliably deduce the position of the true state at any given time is essential to the ability of the bounded controller to “step in” to preserve closed-loop stability in the event of failure of MPC (see Remark 14). One way of ascertaining bounds on the position of the true states, by looking only at the values of the estimates, is described in Proposition 3 below. The proof of this proposition is given in the Appendix.

Proposition 3. Consider the constrained LTI system of Eqs. 1–3, the state observer of Eq. 4 and the bounded control law of Eqs. 12 and 13. Let $T_d^* := (1/\beta) \ln\{\kappa(\beta)e_{\max}(0)/[\epsilon\sqrt{c_{\max}}/\lambda_{\max}(P)]\}$ for some $0 < \epsilon < 1$, where $e_{\max}(0) = \max_{\hat{x}(0), x(0) \in \Omega_b} \|\hat{x}(0) - x(0)\|$. Then, there exists a positive real number $\delta_s^* < c_{\max}$ such that for all $\delta_s \leq \delta_s^*$, and for all $t \geq T_d^*$, $\hat{x}'(t)P\hat{x}(t) \leq \delta_s \Rightarrow x'(t)Px(t) \leq c_{\max}$.

Remark 12. The above proposition establishes the existence of a set, $\Omega_s := \{x \in \mathbb{R}^m : x'Px \leq \delta_s\}$, that allows us to ascertain bounds on the position of the true state, by looking only at whether the estimate is within this set. Specifically, for a given bound on the norm of the estimation error, at any given time, if the state estimate is within Ω_s , then we can conclude with certainty that the true state cannot be outside of Ω , irrespective of the controller being used (see the proof in the Appendix for the mathematical details of this argument).

Hybrid Predictive Control

Although the bounded controller possesses a well-defined region of initial conditions that guarantee closed-loop stability in the presence of constraints, the performance of this controller may not be optimal with respect to a given performance criterion. On the other hand, the model predictive controller is well suited for handling constraints within an optimal control setting; however, the analytical characterization of its set of initial conditions, for which closed-loop stability is guaranteed, is a more difficult task than it is through bounded control. The lack of state measurements introduces an additional layer of complexity in implementing the controllers designed under the assumption of state feedback. In this section, we show how to reconcile the two control approaches by means of a hybrid control structure that combines the desirable properties of both approaches.

To this end, we consider the LTI system of Eqs. 1–3, for which the bounded controller of Eqs. 12 and 13, the state estimator of Eq. 4, and the model predictive controller of Eqs. 9 and 10 have been designed. The control problem is formulated as the one of designing a set of switching laws that orchestrate the transition between the predictive controller and the bounded controller under output feedback in a way that: (1) respects input constraints; (2) guarantees asymptotic stability of the origin of the closed-loop system, starting from any initial condition within arbitrarily large compact subsets of the state feedback stability region of the bounded controller; and (3) utilizes the optimality properties of MPC whenever possible (that is, when closed-loop stability is guaranteed). For a precise statement of the problem, we first cast the system of Eqs. 1–3 as a switched linear system of the form

$$\begin{aligned}\dot{x} &= Ax + Bu_{i(t)} \\ y &= Cx \\ \|u_{i(t)}\| &\leq u_{\max} \\ i(t) &\in \{1, 2\}\end{aligned}\quad (17)$$

where $i: [0, \infty) \rightarrow \{1, 2\}$ is the switching signal, which is assumed to be a piecewise continuous (from the right) function of time, implying that only a finite number of switches is allowed on any finite interval of time. The index $i(t)$, which takes values in the finite set $\{1, 2\}$, represents a discrete state that indexes the control input, $u(\cdot)$, with the understanding that




Figure 3. Hierarchical hybrid predictive control structure merging the bounded and model predictive controllers and a state observer.

$i(t) = 1$ if and only if $u_i[\hat{x}(t)] = b[\hat{x}(t)]$ (that is, the bounded controller is implemented in the closed-loop system), and $i(t) = 2$ if and only if $u_i[\hat{x}(t)] = M_s[\hat{x}(t)]$ (that is, MPC is implemented in the closed-loop system). We note that while we have already used the subscript i in u to denote the different manipulated inputs, we place the index $i(t)$ as a subscript in u to avoid complicating further the notation and this distinction will be made clear throughout this article. Our goal is to construct a switching law, based on the available state estimates

$$i(t) = \psi[\hat{x}(t), t] \quad (18)$$

that provides the set of switching times that ensure asymptotically stabilizing transitions between the predictive and bounded controllers, in the event that the predictive controller is unable to enforce closed-loop stability for a given initial condition within Ω . In the remainder of this section, we present a switching scheme that addresses this problem. Theorem 1 below summarizes the main result (the proof is given in the Appendix).

Theorem 1. Consider the constrained system of Eq. 17, the bounded controller of Eqs. 12 and 13, the state estimator of Eq. 4, and the MPC law of Eqs. 9 and 10. Let $x(0) \in \Omega_b$, $\hat{x}(0) \in \Omega_b$, $\beta \geq \beta^*$, $\delta_s \leq \delta_s^*$, $\Omega_s(T_d^*) = \{x \in \mathbb{R}^n : x'Px \leq \delta_s(T_d^*)\}$, and $0 < T_d < T_{min} := \min\{t \geq 0 : V[x(t)] = \delta_b, V[x(t)] = c_{max}, u(t) \in \mathcal{U}\}$, where Ω_b and β^* were defined in Proposition 2, and δ_s^* and T_d^* were defined in Proposition 3. Let $T_m \geq \max\{T_\phi, T_d^*\}$ be the earliest time for which $\hat{x}(T_m) \in \Omega_s$ and the MPC law prescribes a feasible solution, $M_s[\hat{x}(T_m)]$. Also, let $T_f > T_m$ be the earliest time for which $\dot{V}[\hat{x}(T_f)] \geq 0$. Then, the following switching rule

$$i(t) = \begin{cases} 1, & 0 \leq t < T_m \\ 2, & T_m \leq t < T_f \\ 1, & t \geq T_f \end{cases} \quad (19)$$

where $i(t) = 1$ if and only if $u_i[\hat{x}(t)] = b[\hat{x}(t)]$ and $i(t) = 2$ if and only if $u_i[\hat{x}(t)] = M_s[\hat{x}(t)]$, asymptotically stabilizes the origin of the closed-loop system.

Remark 13. Figure 3 depicts a schematic representation of the hybrid predictive control structure proposed in Theorem 1. The four main components of this structure include the state estimator, the bounded controller, the predictive controller, and a high-level supervisor that orchestrates the switching between the two controllers.

The implementation of the control strategy can be understood through the following stepwise procedure (see Figure 4):

(1) Given the system of Eq. 1 and the constraints of Eq. 3, design the bounded controller using Eqs. 12 and 13. Given the performance objective, design the predictive controller.

(2) Compute the stability region estimate for the bounded controller under state feedback, Ω , using Eqs. 14 and 15 and, for a choice of the output feedback stability region, $\Omega_b \subset \Omega$, compute e_m , β , and T_d defined in Propositions 1 and 2.

(3) Compute the region Ω_s and T_d^* (see Proposition 3), which ensure that $\hat{x} \in \Omega_s \Rightarrow x \in \Omega$ for all times greater than T_d^* .

(4) Initialize the closed-loop system at any initial condition, $x(0)$, within Ω_b , under the bounded controller using an initial guess for the state, $\hat{x}(0)$, within Ω_b . Keep the bounded controller active for a period of time, $[0, T_m)$.

(5) At $t = T_m$ (by which time $\hat{x} \in \Omega_s$), test the feasibility of MPC using values of the estimates generated by the state observer. If MPC yields no feasible solution, keep implementing the bounded controller.

(6) If MPC is feasible, disengage the bounded controller from the closed-loop system and implement the predictive controller instead. Keep MPC active for as long as $\hat{x} \in \Omega_s$ and $V(\hat{x})$ continues to decay.

(7) At the first instance that either $V(\hat{x})$ begins to increase under MPC, or MPC runs into infeasibilities, switch back to the bounded controller and implement it for all future times; else keep the predictive controller active.

Remark 14. An important feature that distinguishes the switching logic of Theorem 1 from the state feedback logic in




Figure 4. Evolution of the closed-loop system under the controller switching strategy proposed in Theorem 1.

Note that MPC is not activated (even if feasible) before $\hat{x}$ enters Ω_s .

(El-Farra et al., 2004b) is the fact that, in the output feedback case, MPC implementation does not begin until some period of time, $[0, T_m)$, has elapsed, even if MPC is initially feasible. The rationale for this delay (during which only bounded control is implemented) is the need to ensure that, by the time MPC is implemented, the norm of the estimation error has already decayed to sufficiently low values that would allow the supervisor to switch back to the bounded controller and preserve closed-loop stability in the event that MPC needs to be switched out after its implementation. More specifically, before MPC can be safely activated, the supervisor needs to make sure that: (1) the estimation error has decayed to levels below e_m , and (2) the state x will be contained within Ω at any time that MPC could be switched out after its implementation. The first requirement ensures that the bounded controller, when (and if) reactivated, is able to tolerate the estimation error and still be stabilizing within Ω (see Proposition 1). This requirement also implies that T_m must be greater than T_d , which is the time beyond which the observer design guarantees $\|e\| \leq e_m$ (see Part 1 of the proof of Proposition 2). The second requirement, on the other hand, implies that MPC cannot be safely switched in before the estimate enters Ω_s and the norm of the error is such that $\hat{x} \in \Omega_s \Rightarrow x \in \Omega$ (see Remarks 11 and 12). Recall from Proposition 3 that this property holds for all $t \geq T_d^*$; thus T_m must also be greater than T_d^* , that is, $T_m \geq \max\{T_d, T_d^*\}$. For times greater than T_m , the implementation of MPC can be safely tested without fear of closed-loop instability, because even if $\hat{x}$ tries to escape Ω_s (possibly indicating that x under MPC is about to escape Ω), the switching rule allows the supervisor to switch back immediately to the bounded controller (before $\hat{x}$ leaves Ω_s), thus guaranteeing that at the switching time $\hat{x} \in \Omega_s$, $x \in \Omega$ and $\|e\| \leq e_m$, which together imply that the bounded controller, once switched in, will be asymptotically stabilizing (see the proof of Theorem 1 for the mathematical details of this argument).

Remark 15. The proposed approach does not require or provide any information as to whether the model predictive controller itself is asymptotically stabilizing, starting from any initial condition within Ω_b . In other words, the approach does not turn Ω_b into a stability region for MPC. What the approach does, however, is turn Ω_b into a stability region for the switched closed-loop system. The value of this can be understood in light of the difficulty in obtaining, a priori, an analytical characterization of the set of admissible initial conditions that the model predictive controller can steer to the origin in the presence of input constraints. Given this difficulty, by embedding MPC implementation within Ω_b and using the bounded controller as a fallback controller, the switching scheme of Theorem 1 allows us to safely initialize the closed-loop system anywhere within Ω_b with the guarantee that the bounded controller can always intervene to preserve closed-loop stability in case the MPC is infeasible or destabilizing (due, for example, to a poor choice of the initial condition and/or improper tuning of the horizon length). This safety feature distinguishes the bounded controller from other fallback controllers that could be used, such as PID controllers, which do not provide a priori knowledge of the constrained stability region and thus do not guarantee a safe transition in the case of unstable plants. For these controllers, a safe transition is critically dependent on issues such as whether the operator is able to properly tune the controller on-line, which, if not achieved, can possibly result in

instability and/or performance degradation. The transition from the MPC to the bounded controller, on the other hand, is not fraught with such uncertainty and is always safe because of the a priori knowledge of the stability region. This aspect helps make plant operation safer and smoother.

Remark 16. Note that, once MPC is switched in, if $V(\hat{x})$ continues to decrease monotonically, then the predictive controller will be implemented for all $t \geq T_m$. In this case, the optimality properties of the predictive controller are practically recovered. Note also, that in this approach, the state-feedback predictive controller design is not required to be robust with respect to state measurement errors (see Grimm et al., 2003, for example, when MPC is not robust) because even if it is not robust, closed-loop stability can always be guaranteed by switching back to the bounded controller (within its associated stability region), which provides the desired robustness with respect to the measurement errors.

Remark 17. Even though the result of Theorem 1 was derived using the MPC formulation of Eqs. 9 and 10, it is important to point out that the same result applies to MPC formulations that use different types of stability constraints (see Remark 3). The basic commonality between the formulation of Eqs. 9 and 10 and other stability-handling formulations is the fact that, in all of these formulations, the implementation of MPC depends on whether the optimization problem, subject to the stability constraints, is feasible. This issue of feasibility is therefore accommodated in the switching logic by requiring the supervisor to check the feasibility of MPC in addition to monitoring the evolution of V . On the other hand, when a conventional MPC formulation (that is, one with no stability constraints) is used, potential infeasibility of the optimization problem is no longer a concern, and, therefore, in this case the supervisory switching logic rests only on monitoring the evolution of V .

Remark 18. In conventional MPC formulations (with no stability constraints), the objective function is typically a tuning parameter that can be manipulated, by changing the weighting matrices, to ensure stability. In the hybrid predictive control framework, on the other hand, because closed-loop stability is guaranteed by the bounded controller, the objective function and the weights, instead of being arbitrarily tuned to guarantee stability, can be chosen to reflect actual physical costs and satisfy desirable performance objectives (for example, fast closed-loop response).

Remark 19. The switching strategy of Theorem 1 can be generalized to allow for multiple switchings between MPC and the bounded controller. According to Theorem 1, once MPC is terminated, the bounded controller is implemented for all future times. In a more general scheme, however, the bounded controller need not stay in the closed-loop system for all future times. Instead, it can be used only until it brings the closed-loop state trajectory sufficiently closer to the origin (for example, to a predetermined inner level set of V) at which point MPC is reactivated (if feasible) and its behavior monitored once again. Potential chattering problems attributed to undesirable back-and-forth switching, are eliminated by allowing only a finite number of switches between the two controllers over any finite time interval. This scheme offers the possibility of further enhancement in the closed-loop performance (over that obtained from the bounded controller) by implementing MPC for

longer periods of time (as permitted by stability considerations).

Remark 20. The hybrid predictive control structure proposed in this article can be further extended to deal with the case when both input and state constraints are present. In one possible extension, state constraints would be incorporated directly as part of the constrained optimization problem that yields the model predictive control law. In addition, an estimate of the stability region for the bounded controller would be obtained by intersecting the region described by Eq. 14 with the region described by the state constraints (presumed closed, convex, and containing the origin in its interior) and computing the largest invariant subset within the intersection. Using this stability region (which now accounts for both input and state constraints) in place of Ω , the implementation of the hybrid predictive control strategy can proceed following the same methodology outlined in Theorem 1.

Remark 21. The proposed switching scheme differs, both in its objective and implementation, from other MPC formulations that involve switching. For example, in dual-mode MPC (Michalska and Mayne, 1993), the strategy involves switching from MPC to a locally stabilizing controller once the state is brought near the origin by MPC. The purpose of switching in this approach is to relax the terminal equality constraint whose implementation is computationally burdensome for nonlinear systems. However, the set of initial conditions for which MPC is guaranteed to steer the state close to the origin is not explicitly known a priori. In contrast, switching from MPC to the bounded controller is used in our work only to prevent any potential closed-loop instability arising from implementing MPC without the a priori knowledge of the admissible initial conditions. Therefore, depending on the stability properties of the chosen MPC, switching may or may not occur, and if it occurs, it can take place near or far from the origin.

Simulation Studies

In this section, two simulation studies are presented to demonstrate the implementation, and evaluate the effectiveness of the proposed hybrid predictive control strategy as well as test its robustness with respect to modeling errors and measurement noise.

Illustrative example

In this section, we demonstrate an application of the proposed hybrid predictive control strategy to a three-dimensional linear system where only two of the states are measured. Specifically, we consider an exponentially unstable linear system of the form of Eqs. 1 and 2 with

$$A = \begin{bmatrix} 0.55 & 0.15 & 0.05 \\ 0.15 & 0.40 & 0.20 \\ 0.10 & 0.15 & 0.45 \end{bmatrix} \quad B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 1 & 1 \end{bmatrix}$$

$$C = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

where both inputs u_1 and u_2 are constrained in the interval $[-1, 1]$ (throughout the text of the example, and the associated figures, subscripts 1 and 2 in u_1 and u_2 refer to the two different

manipulated inputs and do not refer to the values of the discrete variable $i(t)$). Using Eqs. 12 and 13, we initially designed the bounded controller and constructed its stability region Ω using Eqs. 14 and 15. The matrix P was chosen as

$$P = \begin{bmatrix} 6.5843 & 4.2398 & -3.830 \\ 4.2398 & 3.6091 & -2.667 \\ -3.830 & -2.667 & 2.8033 \end{bmatrix}$$

and the observer gain parameter was chosen to be $\beta = 500$ to ensure closed-loop stability for all initial conditions within a set $\Omega_b \subset \Omega$. For the model predictive controller, the parameters in the objective function of Eq. 9 were chosen as $Q = qI$, with $q = 1$ and $R = rI$, with $r = 0.1$. We also chose a horizon length of $T = 1.5$ in implementing the predictive controller of Eqs. 9 and 10. The resulting quadratic program was solved using the MATLAB (MathWorks, Natick, MA) subroutine QuadProg, and the set of ODEs integrated using the MATLAB solver ODE45.

In the first simulation run (solid lines in Figures 5 and 6), the states were initialized at $x_0 = [0.75 \ -0.5 \ 1.0]'$, whereas the observer states were initialized at $\hat{x}_0 = [0 \ 0 \ 0]'$ (which belong to the stability region of the bounded controller, Ω_b). The supervisor uses the bounded controller, while continuously checking the feasibility of MPC. At $t = 5.45$, MPC (based on the state estimates) becomes feasible and is implemented in the closed-loop system to achieve asymptotic stability. The total cost of stabilization achieved by the switching strategy (as measured by Eq. 9) is $J = 80.12$. The switching scheme leads to better performance when compared with the scenario where the bounded controller is implemented for all times (the cost in this case is $J = 121.07$), and also better than the model predictive controller, because, starting from this initial condition, the predictive controller prescribes no feasible control move (equivalent to an infinite value for the objective function). Note that even though feasibility of MPC could have been achieved earlier by increasing the horizon length to $T = 3.5$ (dashed lines in Figures 5 and 6), this conclusion could not be reached a priori, that is, before running the closed-loop simulation in its entirety to check whether the choice $T = 3.5$ is appropriate. In the absence of any a priori knowledge of the necessary horizon length, for a given initial condition, the proposed hybrid predictive control can be used to guarantee closed-loop stability regardless of the choice of the horizon length.

In the second set of simulation results, we demonstrate the need for a choice of an observer gain consistent with the choice of Ω_b . To this end, we consider an observer design with a relatively low gain ($\beta = 0.5$) placing the observer poles at -0.5 , -1.0 , and -1.5 . With the low observer gain, the estimates take a long time to converge to the true state values, resulting in the implementation of "incorrect" control action for a large period of time, by the end of which the states have already escaped out of Ω (even though the states and state estimates were initialized within Ω_b), thus resulting in closed-loop instability (dashed lines in Figures 7 and 8). Recall that stability was achieved when the observer gain was chosen to be $\beta = 500$ in the first set of simulation runs (the state and input profiles for this case are reproduced by the solid lines in Figures 7 and 8 for convenience of comparison).




Figure 5. States of the closed-loop system under the proposed hybrid predictive control strategy using two different horizon lengths of $T = 1.5$ (solid) and $T = 3.5$ (dashed) in MPC implementation, and an observer gain parameter of $\beta = 500$.

To recover, as closely as possible, the state feedback stability region, large values of the observer gain are needed. However, it is well known that high observer gains can amplify measurement noise and induce poor performance. These observations point to a fundamental trade-off that cannot be resolved by simply changing the estimation scheme (see Remark 10). For example, if the observer gain consistent with the choice of the output feedback stability region is abandoned, the noise prob-

lem may disappear, but then stability cannot be guaranteed. One approach to avoid this problem in practice is to initially use a large observer gain that ensures quick decay of the initial estimation error, and then switch to a low observer gain. In the following simulation, we demonstrate how this idea, in conjunction with switching between the controllers, can be used to mitigate the undesirable effects of measurement noise. To illustrate this point, we switch between the high and low observer gains used in the first two simulation runs and demonstrate the attenuation of noise. Specifically, we consider the nominal system described by Eqs. 1 and 2 (see the first paragraph of this subsection for the values of A , B , and C), together with model uncertainty and measurement noise. The model matrix A_m (used for controller and observer design) is assumed to be within 5% error of the system matrix A and the sensors are assumed to introduce noise in the measured outputs as $y(t) = Cx(t) + \delta(t)$, where $\delta(t)$ is a random Gaussian noise with zero mean and a variance of 0.01. As seen by the solid lines in Figure 9, starting from the initial condition, $x_0 = [0.75 \ -0.5 \ 1.0]'$, using a high observer gain followed by a switch to a low observer gain at $t = 1.0$, and a switch from bounded control to MPC at $t = 3.5$, the supervisor is still able to preserve closed-loop stability, while at the same time resulting in a smoother control action (solid lines in Figure 10) when compared to the case where a high




Figure 6. Manipulated input profiles under the proposed hybrid predictive control strategy using two different horizon lengths of $T = 1.5$ (solid) and $T = 3.5$ (dashed) in MPC implementation, and an observer gain parameter $\beta = 500$.




Figure 7. States of the closed-loop system under the proposed hybrid predictive control strategy using two different gain parameters of $\beta = 500$ (solid) and $\beta = 0.5$ (dashed) in observer implementation, and horizon length of $T = 1.5$ in the predictive controller.

observer gain is used for the entire duration of the simulation (dotted lines in Figures 9 and 10).

Application to a chemical reactor example

In this subsection, we consider a nonlinear chemical reactor example to test the robustness of the proposed hybrid predictive control structure (designed on the basis of the linearization of the nonlinear model around the operating steady state) with

respect to modeling errors arising due to the presence of nonlinear terms. Specifically, we consider a well-mixed continuous stirred tank reactor where an irreversible elementary exothermic reaction of the form $A \rightarrow_k B$, takes place, where A is the reactant species and B is the product. The feed to the reactor consists of pure A at flow rate F , molar concentration C_{A0} , and temperature T_{A0} . A jacket is used to remove/provide heat to the reactor. Under standard modeling assumptions, a mathematical model of the process can be derived from material and energy balances and takes the following form

$$\begin{aligned} V \frac{dT_R}{dt} &= F(T_{A0} - T_R) + \frac{(-\Delta H)}{\rho c_p} k_0 \exp\left(\frac{-E}{RT_R}\right) C_A V + \frac{Q}{\rho c_p} \\ V \frac{dC_A}{dt} &= F(C_{A0} - C_A) - k_0 \exp\left(\frac{-E}{RT_R}\right) C_A V \end{aligned} \quad (20)$$

where C_A denotes the concentrations of the species A ; T_R denotes the temperature of the reactor; Q denotes rate of heat input/removal from the reactor; V denotes the volume of the reactor; ΔH , k , and E denote the enthalpy, preexponential constant, and activation energy of the reaction, respectively; and c_p and ρ denote the heat capacity and density of the reactor,




Figure 8. Manipulated input profiles under the proposed hybrid predictive control strategy using two different gain parameters of $\beta = 500$ (solid) and $\beta = 0.5$ (dashed) in observer implementation, and horizon length of $T = 1.5$ in the predictive controller.




Figure 9. States of the closed-loop system under the proposed hybrid predictive control strategy in the presence of measurement noise and modeling errors.

The solid profiles depict the case when the state observer is initially implemented using $\beta = 500$ and then switched to $\beta = 0.5$ at $t = 1.0$. The dotted profiles depict the case when the state observer is implemented using $\beta = 500$ for all times. The horizon length for the predictive controller is $T = 1.5$.

respectively. The values of the process parameters are given in Table 1. It was verified that these conditions correspond to three steady states, one unstable (given in Table 1) and two locally asymptotically stable.

The control problem is to regulate both the outlet concentration of the reactant C_A , and the reactor temperature T_R , at the




Figure 10. Manipulated input profiles under the proposed hybrid predictive control strategy in the presence of measurement noise and modeling errors.

The solid profiles depict the case when the state observer is initially implemented using $\beta = 500$ and then switched to $\beta = 0.5$ at $t = 1.0$. The dotted profiles depict the case when the state observer is implemented using $\beta = 500$ for all times. The horizon length for the model predictive controller is $T = 1.5$.

unstable steady state by manipulating the inlet reactant concentration C_{A0} , and the rate of heat input Q , provided by the jacket. The control objective is to be accomplished in the presence of constraints on the manipulated inputs ($|Q| \leq 1$ kJ/min, $|\Delta C_{A0}| = |C_{A0} - C_{A0s}| \leq 1$ mol/L) using only measurements of the reactor temperature. For the purpose of im-

Table 1. Process Parameters and Steady-State Values

Parameter	Value	Unit of Measure
V	0.1	m^3
R	8.314	$\text{kJ/kmol} \cdot \text{K}$
C_{A0s}	1.0	kmol/m^3
T_{A0s}	310.0	K
ΔH	-4.78×10^4	kJ/kmol
k_0	1.2×10^9	s^{-1}
E	8.314×10^4	kJ/kmol
c_p	0.239	$\text{kJ/kg} \cdot \text{K}$
ρ	1000.0	kg/m^3
F	1.67×10^{-3}	m^3/s
T_{Rs}	395.33	K
C_{As}	0.57	kmol/m^3




Figure 11. Stability region estimate Ω' , based on the linearization of the nonlinear model of Eq. 20.

The solid line depicts the closed-loop state trajectory starting from an initial condition for which MPC (designed based on the linear model with $T = 2$) is feasible and stabilizes the nonlinear closed-loop system, whereas the dotted line depicts the state trajectory starting from an initial condition for which MPC is infeasible, and the hybrid predictive control strategy is implemented instead to enforce closed-loop stability.

plementing the control scheme, the process model of Eq. 20 was linearized around the unstable steady state, and the resulting linear model was used for the design of the hybrid predictive control structure. Using Eqs. 12 and 13, we initially designed the bounded controller and constructed its stability region using Eqs. 14 and 15. The matrix P was chosen as

$$P = \begin{bmatrix} 0.027 & 0.47 \\ 0.47 & 15.13 \end{bmatrix}$$

and the observer gain parameter was chosen to be $\beta = 5$. It should be noted here that, because of the linearization effect, the largest invariant region, Ω , computed using Eqs. 13–15 applied to the linear model, includes physically meaningless initial conditions ($C_A < 0$) and, therefore, a smaller level set, $\Omega' \subset \Omega$, that includes only physically meaningful initial conditions, was chosen and used as the stability region estimate for the bounded controller (see Figure 11). For the predictive controller, the parameters in the objective function of Eq. 9 were chosen as $Q = qI$, with $q = 1$ and $R = rI$, with $r = 1$. We also chose a horizon length of $T = 2$ in implementing the predictive controller of Eqs. 9 and 10. The resulting quadratic program was solved using the MATLAB subroutine QuadProg, and the set of ODEs integrated using the MATLAB solver ODE45. In all simulations, the controllers, designed on the basis of the linearization around the unstable steady state, were implemented on the nonlinear system of Eq. 20.

As seen by the solid lines in Figures 12 and 13, starting from the point $(T_R, C_A) = (395.53 \text{ K}, 0.47 \text{ kmol/m}^3)$, the model predictive controller is able to stabilize the process at the desired steady state $(T_{RS}, C_{AS}) = (395.33 \text{ K}, 0.57 \text{ kmol/m}^3)$. From the initial condition, $(T_R, C_A) = (397.33 \text{ K}, 0.67 \text{ kmol/m}^3)$; however, the model predictive controller yields an infea-

sible solution. Therefore, using the proposed hybrid control strategy (dotted lines in Figures 12 and 13), the supervisor implements the bounded controller in the closed-loop system, while continuously checking the feasibility of MPC. At $t = 0.9 \text{ s}$, MPC is found to be feasible and the supervisor switches to the model predictive controller, which in turn stabilizes the nonlinear closed-loop system at the desired steady state. These results clearly demonstrate a certain degree of robustness that the proposed hybrid predictive control structure possesses with respect to modeling errors (arising in this case because of neglecting the nonlinearities).

Remark 22. In both simulation studies presented above, we have treated robustness as a practical implementation issue and performed simulations that test the performance of the control strategy in the event of some plant-model mismatch to show that the hybrid predictive controller does possess some robustness margins that make it suited for practical implementation. We note that our objective here is not to incorporate robustness explicitly into the controller design, but rather to investigate the




Figure 12. Closed-loop reactant concentration (top) and reactor temperature (bottom) starting from an initial condition for which MPC (designed based on the linear model with $T = 2$) is feasible and stabilizes the nonlinear closed-loop system (solid), and starting from an initial condition for which MPC is infeasible, and the hybrid predictive control strategy is implemented instead to enforce closed-loop stability (dotted).




Figure 13. Rate of heat input (top) and inlet reactant concentration in deviation variable form (bottom) when MPC is implemented (solid) and when the hybrid predictive control strategy is implemented (dotted).

robustness margins of the controllers. The issue of incorporating robustness explicitly into the hybrid predictive control strategy is beyond the scope of the current work and is the subject of other research work (see El-Farra et al., 2003; Mhaskar et al., 2003). Furthermore, within predictive control, a common approach to the robust controller design problem is through the min-max type formulations. However, these formulations are well known to be computationally demanding and, more important, suffer from similar difficulties in identifying the set of initial conditions from which robust closed-loop stability can be guaranteed a priori.

Conclusions

In this work, a hybrid predictive control structure, which seamlessly unites MPC and bounded control, was proposed for the output feedback stabilization of linear systems with input constraints. The hybrid control structure consists of four components, including a bounded controller, a model predictive controller, a state estimator, and a higher-level supervisor that orchestrates the switching between the two controllers. The output feedback bounded controller provides a well-defined stability region (but no optimal performance), whereas the model predictive controller allows for constraint handling within an optimal control setting (but no explicit characteriza-

tion of the stability region). Switching laws were derived to reconcile the two approaches in a way that guarantees: (1) closed-loop stability for all initial conditions within the stability region of the bounded controller, and (2) achievement of the MPC performance whenever the pertinent stability criteria are met. The hybrid predictive control structure was shown to provide a safety net for the practical implementation of output feedback MPC by providing, through off-line computations, a priori knowledge of a large set of initial conditions for which stability of the switched closed-loop system is guaranteed. Finally, simulation studies were presented to demonstrate the implementation and evaluate the effectiveness of the proposed hybrid predictive control strategy, as well as test its robustness with respect to modeling errors and measurement noise.

Acknowledgments

Financial support from the National Science Foundation (Grant CTS-0129571) is gratefully acknowledged.

Literature Cited

- Aufderheide, B., V. Prasad, and B. W. Bequette, "A Comparison of Fundamental Model-Based and Multiple Model Predictive Control," *Proc. 40th IEEE Conf. on Decision and Control*, 4863–4868, Orlando, FL (2001).
- Banerjee, A., and Y. Arkun, "Model Predictive Control of Plant Transitions Using a New Identification Technique for Interpolating Nonlinear Models," *J. Process Control*, **8**, 441 (1998).
- Bitmead, R. R., M. Gevers, and V. Wertz, *Adaptive Optimal Control—The Thinking Man's GPC*, Prentice-Hall, Englewood Cliffs, NJ (1990).
- Chen, H., and F. Allgöwer, "A Quasi-Infinite Horizon Nonlinear Model Predictive Control Scheme with Guaranteed Stability," *Automatica*, **34**, 1205 (1998).
- Chmielewski, D., and V. Manousiouthakis, "On Constrained Infinite-Time Linear Quadratic Optimal Control," *Syst. Control Lett.*, **29**, 121 (1996).
- Christofides, P. D., "Robust Output Feedback Control of Nonlinear Singularly Perturbed Systems," *Automatica*, **36**, 45 (2000).
- Christofides, P. D., and A. R. Teel, "Singular Perturbations and Input-to-State Stability," *IEEE Trans. Autom. Contr.*, **41**, 1645, (1996).
- De Oliveira, N. M. C., and L. T. Biegler, "Constraint Handling and Stability Properties of Model-Predictive Control," *AIChE J.*, **40**, 1138 (1994).
- El-Farra, N. H., and P. D. Christofides, "Integrating Robustness, Optimality and Constraints in Control of Nonlinear Processes," *Chem. Eng. Sci.*, **56**, 1841 (2001a).
- El-Farra, N. H., and P. D. Christofides, "Robust Near-Optimal Output Feedback Control of Nonlinear Systems," *Int. J. Contr.*, **74**, 133, (2001b).
- El-Farra, N. H., and P. D. Christofides, "Bounded Robust Control of Constrained Multivariable Nonlinear Processes," *Chem. Eng. Sci.*, **58**, 325 (2003a).
- El-Farra, N. H., and P. D. Christofides, "Coordinating Feedback and Switching for Control of Hybrid Nonlinear Processes," *AIChE J.*, **49**, 2079 (2003b).
- El-Farra, N. H., P. Mhaskar, and P. D. Christofides, "Hybrid Control of Uncertain Process Systems," *Proc. 8th Int. Symp. on Process Systems and Engineering*, Kuming, China, 814 (2003).
- El-Farra, N. H., P. Mhaskar, and P. D. Christofides, "Hybrid Predictive Control of Nonlinear Systems: Method and Applications to Chemical Processes," *Int. J. Robust & Nonlinear Control*, **14**, 199 (2004a).
- El-Farra, N. H., P. Mhaskar, and P. D. Christofides, "Uniting Bounded Control and MPC for Stabilization of Constrained Linear Systems," *Automatica*, **40**, 101 (2004b).
- Garcia, C. E., D. M. Prett, and M. Morari, "Model Predictive Control: Theory and Practice—A Survey," *Automatica*, **25**, 335 (1989).
- Genceli, H., and M. Nikolaou, "Robust Stability Analysis of Constrained $l(1)$ -Norm Model Predictive Control," *AIChE J.*, **39**, 1954 (1993).
- Grimm, G., M. Messina, A. Teel, and S. Tuna, "Nominally Robust Model Predictive Control with State Constraints," *Proc. 42nd IEEE Conf. on Decision & Control*, Maui, HI, 1413 (2003).
- Hu, T., and Z. Lin, "Composite Quadratic Lyapunov Functions for Con-

strained Control Systems," *Proc. 41st IEEE Conf. on Decision & Control*, 3500–3505, Las Vegas, NV (2002).

Keerthi, S., and E. Gilbert, "Optimal, Infinite Horizon Feedback Laws for a General Class of Constrained Discrete Time Systems," *J. Optimization Theory & Appl.*, **57**, 265 (1988).

Khalil, H. K. *Nonlinear Systems*, 2nd Edition, Macmillan, New York (1996).

Kokotovic, P., and Arcak, M. "Constructive Nonlinear Control: A Historical Perspective," *Automatica*, **37**, 637 (2001).

Kothare, S. L. D., and M. Morari, "Contractive Model Predictive Control for Constrained Nonlinear Systems," *IEEE Trans. Automat. Contr.*, **45**, 1053 (2000).

Lin, Y., and E. D. Sontag, "A Universal Formula for Stabilization with Bounded Controls," *Syst. & Contr. Lett.*, **16**, 393 (1991).

Magni, L., G. D. Nicolao, and R. Scattolini, "A Stabilizing Model-Based Predictive Control Algorithm for Nonlinear Systems," *Automatica*, **37**, 1351 (2001).

Mahmoud, N. A., and H. K. Khalil, "Asymptotic Regulation of Minimum Phase Nonlinear Systems Using Output Feedback," *IEEE Trans. Automat. Contr.*, **41**, 1402 (1996).

Mayne, D. "Nonlinear Model Predictive Control: An Assessment," *Proc. 5th Int. Conf. on Chemical Process Control*, J. C. Kantor, C. E. Garcia, and B. Carnahan, eds., 217–231, CACHE, AIChE (1997).

Mayne, D. Q., J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert, "Constrained Model Predictive Control: Stability and Optimality," *Automatica*, **36**, 789 (2000).

Mhaskar, P., N. H. El-Farra, and P. D. Christofides, "Robust Hybrid Predictive Control of Nonlinear Systems," *Automatica*, in press (2004).

Michalska, H., and D. Q. Mayne, "Receding Horizon Control of Nonlinear Systems," *IEEE Trans. Automat. Contr.*, **35**, 814 (1990).

Michalska, H., and D. Q. Mayne, "Robust Receding Horizon Control of Constrained Nonlinear Systems," *IEEE Trans. Automat. Contr.*, **38**, 1623 (1993).

Michalska, H., and D. Q. Mayne, "Moving Horizon Observers and Observer-Based Control," *IEEE Trans. Automat. Contr.*, **40**, 995 (1995).

Pistikopoulos, E. N., V. Dua, N. A. Bozinis, A. Bemporad, and M. Morari, "On-line Optimization via Off-line Parametric Optimization Tools," *Comput. & Chem. Eng.*, **26**, 175 (2002).

Polak, E., and T. H. Yang, "Moving Horizon Control of Linear Systems with Input Saturation and Plant Uncertainty—Part 1: Robustness," *Int. J. Contr.*, **58**, 613 (1993).

Primbs, J. A., V. Nevistic, and J. C. Doyle, "A Receding Horizon Generalization of Pointwise Min-Norm Controllers," *IEEE Trans. Automat. Contr.*, **45**, 898 (2000).

Qin, S. J., and T. J. Badgwell, "An Overview of Industrial Model Predictive Control Technology," *Proc. 5th Int. Conf. on Chemical Process Control*, J. C. Kantor, C. E. Garcia, and B. Carnahan, eds., 232–256, CACHE, AIChE (1997).

Rao, C. V., and J. B. Rawlings, "Constrained Process Monitoring: Moving-Horizon Approach," *AIChE J.*, **48**, 97 (2002).

Rawlings, J. B., and K. R. Muske, "The Stability of Constrained Receding Horizon Control," *IEEE Trans. Automat. Contr.*, **38**, 1512 (1993).

Richalet, J., A. Rault, J. L. Testud, and J. Papon, "Model Predictive Heuristic Control: Applications to Industrial Processes," *Automatica*, **14**, 413 (1978).

Rugh, W. J., and J. S. Shamma, "Research on Gain Scheduling," *Automatica*, **36**, 1401 (2000).

Schott, K. D., and B. W. Bequette, "MMAC: Using a Bank of Linear Models to Control a Nonlinear System," *NATO ASI on Nonlinear Model-Based Process Control*, R. Berber and C. Kravaris, eds., 33–57, Kluwer Academic, Dordrecht, The Netherlands (1998).

Shamma, J. S., and K. Tu, "Output Feedback Control for Systems with Constraints and Saturations: Scalar Control Case," *Syst. & Contr. Lett.*, **25**, 1 (1998).

Sun, D., and K. A. Hoo, "Dynamic Transition Control Structure for a Class of SISO Nonlinear Systems," *IEEE Trans. Contr. Syst. Technol.*, **7**, 622 (1999).

Sussmann, H. J., and P. V. Kokotovic, "The Peaking Phenomenon and the Global Stabilization of Nonlinear Systems," *IEEE Trans. Automat. Contr.*, **36**, 424 (1991).

Sznaier, M., R. Suarez, and J. Cloutier, "Suboptimal Control of Constrained Nonlinear Systems via Receding Horizon Constrained Control Lyapunov Functions," *Int. J. Robust & Nonlin. Contr.*, **13**, 247 (2003).

Teel, A., and L. Praly, "Global Stabilizability and Observability Imply

Semi-Global Stabilizability by Output Feedback," *Syst. & Contr. Lett.*, **22**, 313 (1994).

Appendix

Proof of Proposition 1. The proof consists of two parts. In the first part, we establish that the bounded state feedback control law of Eqs. 12 and 13 enforces asymptotic stability for all initial conditions within Ω . In the second part, we compute an estimate of the tolerable measurement error, e_m , which guarantees that a state trajectory starting within Ω remains within it for all $\|e\| \leq e_m$, for all $t \geq 0$.

Part 1. Substituting the state feedback control law of Eqs. 12 and 13 into the system of Eq. 1 and evaluating the time derivative of the Lyapunov function along the closed-loop trajectories, it can be shown that

$$\begin{aligned} \dot{V}(x) &= L_f V(x) + L_g V(x) u(x) \\ &\leq \frac{-\rho x' P x}{[1 + \sqrt{1 + (2u_{\max} \|B' P x\|)^2}]^2} \end{aligned} \quad (A1)$$

for all $x \in \Phi$, and thus for all $x \in \Omega$, where Φ and Ω were defined in Eqs. 14 and 15, respectively. Because the denominator term in Eq. A1 is bounded on Ω , there exists a positive real number ρ^* , such that

$$\dot{V} \leq -\rho^* x' P x \quad (A2)$$

for all $x \in \Omega$, which implies that the origin of the closed-loop system, under the control law of Eqs. 12 and 13, is asymptotically stable, with Ω as an estimate of the domain of attraction.

Part 2. In this part, we analyze the behavior of $\dot{V}$ on the boundary of Ω [that is, the level surface described by $V(x) = c_{\max}$] under bounded measurement errors, $\|e\| \leq e_m$. To this end, we have

$$\begin{aligned} \dot{V}(x) &= L_f V(x) + L_g V(x) u(x - e) \\ &= L_f V(x) + L_g V(x) u(x) + L_g V(x) [u(x - e) - u(x)] \\ &\leq -\rho^* c_{\max} + \|L_g V\| \|u(x - e) - u(x)\| \\ &\leq -\rho^* c_{\max} + M \|u(x - e) - u(x)\| \end{aligned} \quad (A3)$$

for all $x = V^{-1}(c_{\max})$, where $M = \max_{V(x)=c_{\max}} (\|L_g V(x)\|)$ (note that M exists given that $\|L_g V(\cdot)\|$ is continuous and the maximization is considered on a closed set). Because $u(\cdot)$ is continuous, then given any positive real number r such that $\mu = (\rho^* c_{\max} - r)/M > 0$, there exists $e_m > 0$ such that if $\|x - e\| = \|e\| \leq e_m$, then $\|u(x - e) - u(x)\| \leq \mu$ and, consequently

$$\dot{V}(x) \leq -\rho^* c_{\max} + M\mu = -r < 0 \quad (A4)$$

for all $x = V^{-1}(c_{\max})$. This implies that for all measurement errors such that $\|e\| \leq e_m$, we have $\dot{V} < 0$ on the boundary of Ω . Therefore, under the bounded controller, any closed-loop state trajectory, starting within Ω , cannot escape this region, that is, $x(t) \in \Omega \forall t \geq 0$. This completes the proof of the proposition.

Proof of Proposition 2. The proof is divided into two parts. In the first part, we show that given $\delta_b > 0$ (the size of the output feedback stability region) such that $\Omega_b \subset \Omega$, there exists a choice of β such that the closed-loop trajectories are bounded for all $x(0)$ and $\hat{x}(0)$ belonging in Ω_b . In the second part, we use a Lyapunov argument, together with boundedness of the states of the closed-loop system and of the estimation error, to show that the state is ultimately bounded with a bound that depends on the norm of the estimation error, and converges to zero to zero.

Part 1. From Eq. 5, we have that the error dynamics are given by $\dot{e} = (A - LC)e$, where $A - LC$ is Hurwitz, and all the eigenvalues of the matrix $A - LC$ satisfy $\lambda \leq -\beta$. It then follows that an estimate of the form $\|e(t)\| \leq \kappa(\beta) \|e(0)\| \exp(-\beta t)$ holds for some $\kappa(\beta) > 0$, for all $t \geq 0$. Given any positive real number, δ_b , such that $\Omega_b = \{x \in \mathbb{R}^n : \|x\|_P^2 \leq \delta_b\} \subset \Omega$, let $T_{\min} = \min\{t \geq 0 : V[x(0)] = \delta_b, V[x(t)] = c_{\max}, u(t) \in \mathcal{U}\}$ (that is, $T_{\min}$ is the shortest time during which the closed-loop state trajectory can reach the boundary of Ω starting from the boundary of Ω_b using any admissible control action). Furthermore, let $e_{\max}(0) = \max_{x, \hat{x} \in \Omega_b} \|x(0) - \hat{x}(0)\|$ [$e_{\max}(0)$ therefore is the largest possible initial error given that both the states and state estimates are initialized within Ω_b]. Choose T_d such that $0 < T_d < T_{\min}$ and let β^* be such that $e_m \leq \kappa(\beta^*) e_{\max}(0) \exp(-\beta^* T_d)$ [the existence of such a β^* follows from the fact that $\kappa(\beta)$ is polynomial in β]. For any choice of $\beta \geq \beta^*$, therefore, it follows that $\|e(T_{\min})\| \leq e_m$ [given that $\|e(T_d)\| \leq e_m$, $T_{\min} \geq T_d$ and the bound on the norm of the estimation error decreases monotonically with time for all $t \geq 0$]. This implies that the norm of the estimation error decays to a value less than e_m before the closed-loop state trajectory, starting within Ω_b , could reach the boundary of Ω . It then follows from Proposition 1 that the closed-loop state trajectory cannot escape Ω for all $t \geq 0$, that is, the trajectories are bounded and $\|x(t)\|_P^2 \leq c_{\max} \forall t \geq 0$.

Part 2. To prove asymptotic stability, we note that for all $x \in \Omega$

$$\begin{aligned} \dot{V}(x) &= L_f V(x) + L_g V(x) u(x - e) \\ &= L_f V(x) + L_g V(x) u(x) + \|L_g V(x)\| \|u(x - e) - u(x)\| \\ &\leq -\rho^* \|x\|_P^2 + M \|u(x - e) - u(x)\| \end{aligned} \quad (A5)$$

The term $\|u(x - e) - u(x)\|$ is continuous and vanishes when $e = 0$. Therefore, given that both x and e are bounded, there exists a positive real number $\varphi(e_m)$ such that $\|u(x - e) - u(x)\| \leq \varphi(e_m)$ for all $\|x\|_P^2 \leq c_{\max}$, $\|e\| \leq e_m$. Substituting this estimate into Eq. A5 yields

$$\begin{aligned} \dot{V}(x) &\leq -\rho^* \|x\|_P^2 + M \varphi(e_m) \leq -\frac{\rho^*}{2} \|x\|_P^2 \\ \forall \|x\|_P &\geq \sqrt{\frac{2M\varphi(e_m)}{\rho^*}} := \gamma_1(\|e\|) \end{aligned} \quad (A6)$$

where $\gamma_1(\cdot)$ is a class $\mathcal{H}$ function [a continuous function $\alpha(\cdot)$ is said to belong to class $\mathcal{H}$ if it is strictly increasing and $\alpha(0) = 0$; see also Khalil, 1996]. The above inequality implies that $\dot{V}$ is negative outside some residual set whose size depends on $\|e\|$. Using the result of Theorem 5.1–Corollary 5.2 in (Khalil,

1996), this implies that, for any $x(0) \in \Omega_b$, there exists a class $\mathcal{KL}$ function $\beta(\cdot, \cdot)$ and a class $\mathcal{K}$ function $\gamma_2(\cdot)$, such that

$$\|x(t)\| \leq \beta(\|x(0)\|, t) + \gamma_2(\sup_{\tau \geq 0} \|e(\tau)\|) \forall t < \infty \quad (1)$$

implying that the x subsystem of Eq. 16, with e as input, is input-to-state stable (recall, from Proposition 1, that $x(t) \in \Omega_b \forall t \geq 0$). Noting also that the e subsystem of Eq. 16 is asymptotically stable ($\lim_{t \rightarrow \infty} \|e(t)\| = 0$), and using Lemma 5.6 in (Khalil, 1996), we get that the interconnected system of Eq. 16 is asymptotically stable. This completes the proof of the proposition.

Proof of Proposition 3. Because the error dynamics obey a bound of the form $\|e(t)\| \leq \kappa(\beta) e_{\max}(0) \exp(-\beta t)$, substituting T_d^* into this expression yields $\|e(T_d^*)\| := e^* \leq \sqrt{c_{\max}/\lambda_{\max}(P)}$. Then, for all $t \geq T_d^*$, if $\hat{x}'(t) P \hat{x}(t) \leq \delta_s$ for some $\delta_s > 0$, we can write

$$\begin{aligned} x'(t) P x(t) &= [\hat{x}(t) + e(t)]' P [\hat{x}(t) + e(t)] \\ &= \hat{x}'(t) P \hat{x}(t) + 2\hat{x}'(t) P e(t) + e'(t) P e(t) \\ &\leq \delta_s + 2\|P \hat{x}(t)\| \|e(t)\| + \|e(t)\|_P^2 \\ &\leq \delta_s + 2\sqrt{\frac{\lambda_{\max}(P^2)\delta_s}{\lambda_{\min}(P)}} e^* + \lambda_{\max}(P) e^{*2} \\ &:= f(\delta_s) \end{aligned} \quad (A9)$$

Note that $f(\cdot)$ is a continuous, monotonically increasing function of $\delta_s \geq 0$, with $f(0) = e^{*2} c_{\max} < c_{\max}$ and $f(c_{\max}) > c_{\max}$. This implies that there exists $0 < \delta_s^* < c_{\max}$ such that, for all $\delta_s \leq \delta_s^*$, the relation $f(\delta_s) \leq c_{\max}$ holds, that is, $x'(t) P x(t) \leq c_{\max}$. This completes the proof of the proposition.

Proof of Theorem 1. Given $x(0) \in \Omega_b$, $\hat{x}(0) \in \Omega_b$, $\beta \geq \beta^*$ (see Proposition 2), and $\delta_s \leq \delta_s^*$ (see Proposition 3), we consider the following cases.

Case 1. Consider first the case when $T_m = \infty$ (that is, MPC is never feasible). Then, we have from Eq. 19 that $i(t) = 1$ and $u[\hat{x}(t)] = b[\hat{x}(t)]$ for all $t \geq 0$. It then follows from Proposition 2 that the origin of the closed-loop system is asymptotically stable.

Case 2. Now, consider the case when $T_m < \infty$ and $T_f = \infty$. Given that $T_m \geq T_d$ and $\|e(T_d)\| \leq e_m$ (see Part 1 of the proof of Proposition 2), we have that $\|e(T_m)\| \leq e_m$. Because only the bounded controller is implemented, that is, $i(t) = 1$, for $0 \leq t < T_m$, it follows from Proposition 1 that $x(t) \in \Omega$ for all $0 \leq t < T_m$ [or that $x(T_m^-) \in \Omega$]. This fact, together with the continuity of the solution of the switched closed-loop system, which follows from the fact that the righthand side of Eq. 17 is continuous in x and piecewise continuous in time—given that only a finite number of switches is allowed over any finite time interval—implies that, upon switching (instantaneously) to the model predictive controller at $t = T_m$, we have $x(T_m) \in \Omega$. Because $T_f = \infty$, then from the definition of T_f in Theorem 1 it follows that $u[\hat{x}(t)] = M_s[\hat{x}(t)]$ and $\dot{V}[\hat{x}(t)] < 0$ for all $t \geq T_m$. This implies that $\hat{x}(t) \in \Omega_s$ for all $t \geq T_m$ and, consequently from the definition of T_m (note that $T_m \geq T_d^*$, where T_d^* was defined in Proposition 3), that $x(t) \in \Omega$ for all $t \geq T_m$ (that is, the closed-loop state trajectory is bounded). Furthermore, given that $\hat{x}(t) = x(t) - e(t)$ and $\lim_{t \rightarrow \infty} \|e(t)\| = 0$, it follows that $\lim_{t \rightarrow \infty} \hat{x}'(t) P \hat{x}(t) = 0$; and thus $\lim_{t \rightarrow \infty} x'(t) P x(t) = 0$. The

origin of the switched closed-loop system is therefore asymptotically stable.

Case 3. Finally, consider the case when $T_m < \infty$ and $T_f < \infty$. From the analysis in case 2 above, we have that $x(t) \in \Omega \forall 0 \leq t < T_f$ [or that $x(T_f^-) \in \Omega$. This fact, together with the continuity of the solution of the switched closed-loop system, implies that, upon switching (instantaneously) to the bounded controller at $t = T_f$, we have $x(T_f) \in \Omega$ and $u(t) = b[\hat{x}(t)]$ for all $t \geq T_f$. We also have $\|e(t)\| \leq e_m$ for all $t \geq T_f$ given that T_f

$> T_d$. Therefore, from Proposition 1 it is guaranteed that $x(t) \in \Omega$ for all $t \geq T_f$. Finally, following the same arguments presented in part 2 of the proof of Proposition 2, it can be shown that $\|x(t)\| \rightarrow 0$ as $t \rightarrow \infty$, which together with boundedness of the state trajectory, establishes asymptotic stability of the origin of the closed-loop system. This completes the proof of the theorem.

Manuscript received Apr. 25, 2003, and revision received Sept. 26, 2003.