Nonlinear Output-Feedback Model Predictive Control with Moving Horizon Estimation

David A. Copp and João P. Hespanha

Abstract— **We introduce an output-feedback approach to model predictive control that combines state estimation and control into a single min-max optimization. Under appropriate assumptions that ensure controllability and observability of the nonlinear process to be controlled, we prove that the state of the system remains bounded and establish bounds on the tracking error for trajectory tracking problems. The results apply both to infinite and finite-horizon optimizations, the latter requiring reversible dynamics and the use of a terminal cost that is an ISS-control Lyapunov function with respect to a disturbance input. A numerical example is presented that illustrates these results.**

I. Introduction

Advances in computer technology have made online optimization a viable and powerful tool for solving control problems in practical applications. Model predictive control (MPC) is an approach that uses online optimization to solve an open-loop optimal control problem at each sampling time and is now quite mature, as evidenced by [1–3]. Several papers in this area are focused on the robustness to model uncertainty, input disturbances, and measurement noise. These studies include robust, worst-case, and min-max MPC. Robust and worst-case MPC are discussed in works such as [4–7]. Min-max MPC for constrained linear systems is considered in [8, 9], and a game theoretic approach for robust constrained nonlinear MPC is proposed in [10]. Nominal or inherent robustness of MPC has also been studied in [3, 11].

MPC is often formulated assuming full-state feedback. In practical cases, however, the full state often cannot be measured and is not available for feedback. This motivates the investigation of output-feedback MPC in which an independent algorithm for state estimation is often used. Examples of algorithms for state estimation include observers, filters, and moving horizon estimation, some of which are discussed in [12]. Of these methods, moving horizon estimation (MHE) is attractive for use with MPC because it explicitly handles constraints and computes the optimal current estimate of the state by solving an online optimization problem over a fixed number of past measurements. Therefore, the computational cost does not grow as more measurements become available. Nonlinear MPC and MHE are both discussed in [13]. A useful overview of constrained nonlinear moving horizon state estimation is given in [14], and more recent results regarding stability of MHE can be found in [15].

D. A. Copp and J. P. Hespanha are with the Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, CA 93106 USA. dacopp@engr.ucsb.edu,hespanha@ece.ucsb.edu

Thus far, results on the stability of output-feedback control schemes based on MPC and MHE (especially for nonlinear systems) are limited. Some joint stability results for state estimation and control of linear systems are given in [16], but output-feedback is not considered. Results on robust output-feedback MPC for constrained linear systems can be found in [17] using a state observer for estimation, and in [18] using MHE for estimation. Fewer results are available for nonlinear output-feedback MPC, although notable exceptions are [3, 19, 20]. Recent studies of input-to-state stability of min-max MPC can be found in [21–23]; however, these references also do not investigate the use of output-feedback.

In this paper, we consider the output-feedback of nonlinear systems with uncertainty and disturbances and formulate the MPC problem as a min-max optimization. In this formulation, a desired cost function is maximized over disturbance and noise variables and minimized over control input variables. In this way, we can solve both the MPC and MHE problems using a single min-max optimization, which gives us an optimal control input sequence at each sampling time for a worst-case estimate of the current state. For both infinite-horizon and finite-horizon optimizations, we show that the state remains bounded under the proposed feedback control law. We also show that the tracking error in trajectory tracking problems is bounded in the presence of measurement noise and input disturbances.

The main assumption for these results is that a saddle-point solution exists for the min-max optimization that is solved at each sampling time. This assumption is a common requirement in game theoretical approaches to control design [24] and presumes observability and controllability of the closed-loop system. For the finite-horizon case, we additionally require that the dynamics are reversible and that there exists a terminal cost that is an ISS-control Lyapunov function with respect to a disturbance input.

The paper is organized as follows. In Section II, we formulate the control problem we would like to solve and discuss its relationship to MPC and MHE. In Section III, we state the main closed-loop stability results. Simulation results showing robust trajectory tracking in the presence of additive disturbances and measurement noise are presented in Section IV. Finally, we provide some conclusions and directions for future research in Section V.

II. Problem Formulation

We consider the control of a time-varying nonlinear discrete-time process of the form

$$x_{t+1} = f_t(x_t, u_t, d_t), \quad y_t = g_t(x_t) + n_t, \quad \forall t \in \mathbb{Z}_{\geqslant 0} \quad (1)$$

with *state* x_t taking values in a set $\mathcal{X} \subset \mathbb{R}^{n_x}$. The inputs to this system are the *control input* u_t that must be restricted to the set $\mathcal{U} \subset \mathbb{R}^{n_u}$, the *unmeasured disturbance* d_t that is known to belong to the set $\mathcal{D} \subset \mathbb{R}^{n_d}$, and the *measurement noise* n_t belonging to the set $\mathcal{N} \subset \mathbb{R}^{n_n}$. The signal y_t, belonging to the set $\mathcal{Y} \subset \mathbb{R}^{n_y}$, denotes the *measured output* that is available for feedback. The *control objective* is to select the control signal $u_t \in \mathcal{U}$, $\forall t \in \mathbb{Z}_{\geqslant 0}$ so as to minimize a criterion of the form

$$\sum_{t=0}^{\infty} c_t(x_t, u_t, d_t) - \sum_{t=0}^{\infty} \eta_t(n_t) - \sum_{t=0}^{\infty} \rho_t(d_t), \tag{2}$$

for worst-case values of the unmeasured disturbance $d_t \in \mathcal{D}$, $\forall t \in \mathbb{Z}_{\geqslant 0}$ and the measurement noise $n_t \in \mathcal{N}$, $\forall t \in \mathbb{Z}_{\geqslant 0}$. The functions $c_t(\cdot)$, $\eta_t(\cdot)$, and $\rho_t(\cdot)$ in (2) are all assumed to take non-negative values. The negative sign in front of $\rho_t(\cdot)$ penalizes the maximizer for using large values of d_t. Boundedness of (2) by a constant γ guarantees that $\sum_{t=0}^{\infty} c_t(x_t, u_t, d_t) \leqslant \gamma + \sum_{t=0}^{\infty} \eta_t(n_t) + \sum_{t=0}^{\infty} \rho_t(d_t)$.

In what follows, we allow the functions $\eta_t(\cdot)$ and $\rho_t(\cdot)$ in the criterion (2) to take the value $+\infty$. This provides a convenient formalism to consider bounded disturbances and noise while formally allowing n_t and d_t to take values in the whole spaces $\mathbb{R}^{n_n}$ and $\mathbb{R}^{n_d}$, respectively. Specifically, considering *extended-value extensions* [25] of the form

$$\rho_t(d_t) \coloneqq \begin{cases} \bar{\rho}_t(d_t) & d_t \in \mathcal{D} \\ \infty & d_t \notin \mathcal{D}, \end{cases} \quad \eta_t(n_t) \coloneqq \begin{cases} \bar{\eta}_t(n_t) & n_t \in \mathcal{N} \\ \infty & n_t \notin \mathcal{N}, \end{cases} \tag{3}$$

with $\bar{\rho}_t$ and $\bar{\eta}_t$ bounded in $\mathcal{D}$ and $\mathcal{N}$, respectively, the minimization of (2) with respect to the control signal u_t need not consider cases where d_t and n_t take values outside $\mathcal{D}$ and $\mathcal{N}$, respectively, as this would directly lead to the cost $-\infty$ for any control signal u_t that keeps the positive term bounded.

Remark 1: While the results presented here are general, the reader is encouraged to consider the quadratic case $c_t(x_t, u_t, d_t) \coloneqq \|x_t\|^2 + \|u_t\|^2$, $\eta_t(n_t) \coloneqq \|n_t\|^2$, $\rho_t(d_t) \coloneqq \|d_t\|^2$ to gain intuition on the results. In this case, boundedness of (2) would guarantee that the state x_t and input u_t are ℓ_2 provided that the disturbance d_t and noise n_t are also ℓ_2. □

A. Infinite-Horizon Online Optimization

To overcome the conservativeness of an open-loop control, we use online optimization to generate the control signals. Specifically, at each time $t \in \mathbb{Z}_{\geqslant 0}$, we compute the control u_t so as to minimize

$$\sum_{s=t}^{\infty} c_s(x_s, u_s, d_s) - \sum_{s=0}^{t} \eta_s(n_s) - \sum_{s=0}^{\infty} \rho_s(d_s) \tag{4}$$

under worst-case assumptions on the *unknown* system's initial condition x_0, unmeasured disturbances d_t, and measurement noise n_t, subject to the constraints imposed by the system dynamics and the measurements y_t collected up to the current time t. Since the goal is to optimize this cost at the current time t to compute the control inputs at times $s \geqslant t$, there is no point in penalizing the running cost $c_s(x_s, u_s, d_s)$ for past time instants $s < t$, which explains the fact that the first summation in (4) starts at time t. There is also no point in considering the values of future measurement noise at times $s > t$, as they will not affect choices made at time t, which explains the fact that the second summation in (4) stops at time t. However, we do need to consider all values for the unmeasured disturbance d_s because past values affect the (unknown) current state x_t, and future values affect the future values of the running cost.

The following notation facilitates formalizing the control law proposed: Given a discrete-time signal $z : \mathbb{Z}_{\geqslant 0} \to \mathbb{R}^n$ and two times $t_0, t \in \mathbb{Z}_{\geqslant 0}$ with $t_0 \leqslant t$, we denote by $z_{t_0:t}$ the sequence $\{z_{t_0}, z_{t_0+1}, ..., z_t\}$. Given a control input sequence $u_{t_0:t-1}$ and a disturbance input sequence $d_{t_0:t-1}$, we denote by[1] $\varphi(t; t_0, x_0, u_{t_0:t-1}, d_{t_0:t-1})$ the state x_t of the system (1) at time t for the given inputs and initial condition $x_{t_0} = x_0$. In addition, to facilitate expressing the corresponding output and running cost, we define

$$\begin{aligned} &g\varphi(t; t_0, x_0, u_{t_0:t-1}, d_{t_0:t-1}) \\ &\qquad\qquad \coloneqq g_t\big(\varphi(t; t_0, x_0, u_{t_0:t-1}, d_{t_0:t-1})\big), \\ &c\varphi(t; t_0, x_0, u_{t_0:t}, d_{t_0:t}) \\ &\qquad\qquad \coloneqq c_t\big(\varphi(t; t_0, x_0, u_{t_0:t-1}, d_{t_0:t-1}), u_t, d_t\big). \end{aligned}$$

This notation allows us to re-write (4) as

$$\begin{aligned} J_t^{\infty}(x_0, u_{0:\infty}, d_{0:\infty}, y_{0:t}) &\coloneqq \sum_{s=t}^{\infty} c\varphi(s; 0, x_0, u_{0:s}, d_{0:s}) \\ - \sum_{s=0}^{t} \eta_s\big(y_s - g\varphi(s; 0, x_0, u_{0:s-1}, d_{0:s-1})\big) &- \sum_{s=0}^{\infty} \rho_s(d_s), \end{aligned} \tag{5}$$

which emphasizes the dependence of (4) on the unknown initial state x_0, the unknown disturbance input sequence $d_{0:\infty}$, the measured output sequence $y_{0:t}$, and the control input sequence $u_{0:\infty}$. Regarding the latter, one should recognize that $u_{0:\infty}$ is composed of two distinct sequences: the (known) past inputs $u_{0:t-1}$ that have already been applied, and the future inputs $u_{t:\infty}$ that still need to be selected.

At a given time $t \in \mathbb{Z}_{\geqslant 0}$, we do not know the value of the variables x_0 and $d_{0:\infty}$ on which the value of criterion (5) depends, so we optimize this criterion under worst-case assumptions on these variables, leading to the following min-max optimization

$$\min_{\hat{u}_{t:\infty|t} \in \mathcal{U}} \max_{\hat{x}_{0|t} \in \mathcal{X}, \hat{d}_{0:\infty|t} \in \mathcal{D}} J_t^{\infty}(\hat{x}_{0|t}, u_{0:t-1}, \hat{u}_{t:\infty|t}, \hat{d}_{0:\infty|t}, y_{0:t}), \tag{6}$$

where the arguments $u_{0:t-1}, \hat{u}_{t:\infty|t}$ to the function $J_t^{\infty}(\cdot)$ in (6) correspond to the argument $u_{0:\infty}$ in the definition of $J_t^{\infty}(\cdot)$ in the left-hand side of (5). The subscript $\cdot_{|t}$ in the (dummy) optimization variables in (6) emphasizes that this optimization is repeated at each time step $t \in \mathbb{Z}_{\geqslant 0}$. At different time steps, these optimizations typically lead

[1] When $t = t_0$, it is understood that we drop all terms that depend on previous values of t, i.e., we write $\varphi(t_0; t_0, x_0)$.

to different solutions, which generally do not coincide with the real control input, disturbances, and noise. We can view the optimization variables $\hat{x}_{0|t}$ and $\hat{d}_{0:\infty|t}$ as (worst-case) estimates of the initial state and disturbances, respectively, based on the past inputs $u_{0:t-1}$ and outputs $y_{0:t}$ available at time t.

Inspired by model predictive control, at each time t, we use as the control input the first element of the sequence

$$\hat{u}^*_{t:\infty|t} = \{\hat{u}^*_{t|t}, \hat{u}^*_{t+1|t}, \hat{u}^*_{t+2|t}, \dots\} \in \mathcal{U}$$

that minimizes (6), leading to the following control law:

$$u_t = \hat{u}^*_{t|t}, \quad \forall t \geqslant 0. \tag{7}$$

B. Finite-Horizon Online Optimization

To avoid solving the infinite-dimensional optimization in (6) that resulted from the infinite-horizon criterion (4), we also consider a finite-horizon version of the criterion (4) of the form

$$\sum_{s=t}^{t+T-1} c_s(x_s, u_s, d_s) + q_{t+T}(x_{t+T}) - \sum_{s=t-L}^{t} \eta_s(n_s) - \sum_{s=t-L}^{t+T-1} \rho_s(d_s), \tag{8}$$

where now the optimization criterion only contains $T \in \mathbb{Z}_{\geqslant 1}$ terms of the running cost $c_s(x_s, u_s, d_s)$, which recede as the current time t advances. The optimization criterion also only contains $L+1 \in \mathbb{Z}_{\geqslant 1}$ terms of the measurement cost $\eta_s(n_s)$. Specifically, the summations in the criterion evaluated at time t, which in (5) started at time 0 and went up to time $+\infty$, now start at time $t-L$ and only go up to time $t+T-1$. We also added a terminal cost $q_{t+T}(x_{t+T})$ to penalize the "final" state at time $t+T$. Defining

$$\begin{aligned} &q\varphi(t; t-L, x_{t-L}, u_{t-L:t-1}, d_{t-L:t-1}) \\ &\qquad \coloneqq q_t\big(\varphi(t; t-L, x_{t-L}, u_{t-L:t-1}, d_{t-L:t-1})\big), \end{aligned}$$

the cost (8) leads to the following finite-dimensional optimization

$$\min_{\hat{u}_{t:t+T-1|t} \in \mathcal{U}} \max_{\hat{x}_{t-L|t} \in \mathcal{X}, \hat{d}_{t-L:t+T-1|t} \in \mathcal{D}} J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}), \tag{9}$$

where

$$\begin{aligned} &J_t(x_{t-L}, u_{t-L:t+T-1}, d_{t-L:t+T-1}, y_{t-L:t}) \\ &\quad \coloneqq \sum_{s=t}^{t+T-1} c\varphi(s; t-L, x_{t-L}, u_{t-L:s}, d_{t-L:s}) \\ &\quad + q\varphi(t+T; t-L, x_{t-L}, u_{t-L:t+T-1}, d_{t-L:t+T-1}) \\ &- \sum_{s=t-L}^{t} \eta_s\big(y_s - g\varphi(s; t-L, x_{t-L}, u_{t-L:s-1}, d_{t-L:s-1})\big) \\ &- \sum_{s=t-L}^{t+T-1} \rho_s(d_s). \end{aligned} \tag{10}$$

In this formulation, we still use a control law of the form (7), but now $\hat{u}^*_{t|t}$ denotes the first element of the sequence $\hat{u}^*_{t:t+T-1|t}$ that minimizes (9).

C. Relationship with Model Predictive Control

When the state of (1) can be measured exactly and the maps $d_t \mapsto f_t(x_t, u_t, d_t)$ are injective (for each fixed x_t and u_t), the initial state x_{t-L} and past values for the disturbance $d_{t-L:t-1}$ are uniquely defined by the "measurements" $x_{t-L:t}$. In this case, the control law (7) that minimizes (9) can also be determined by the optimization

$$\min_{\hat{u}_{t:t+T-1|t} \in \mathcal{U}} \max_{\hat{d}_{t:t+T-1|t} \in \mathcal{D}} J_t(x_{t-L}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, d_{t-L:t-1}, \hat{d}_{t:t+T-1|t}),$$

with

$$\begin{aligned} J_t(x_t, u_{t:t+T-1}, d_{t:t+T-1}) &\coloneqq \sum_{s=t}^{t+T-1} c\varphi(s; t, x_t, u_{t:s}, d_{t:s}) \\ &+ q\varphi(t+T; t, x_t, u_{t:t+T-1}, d_{t:t+T-1}) - \sum_{s=t}^{t+T-1} \rho_s(d_s), \end{aligned}$$

which is essentially the robust model predictive control problem with terminal cost considered in [10, 26].

Remark 2 (Economic MPC): It is worth noting that our framework is more general than standard forms of MPC where the minimal cost is achieved at the optimal feasible state and input in order to ensure stability of the desired state. It can also apply to economic MPC where the operating cost of the plant is used directly in the objective function, and therefore the cost need not be zero or minimal at the optimal state and input [27]. □

D. Relationship with Moving-Horizon Estimation

When setting both $c_s(\cdot)$ and $q_{t+T}(\cdot)$ equal to zero in the criterion (10), this optimization no longer depends on $u_{t:t+T-1}$ and $d_{t:t+T-1}$, so the optimization in (9) simply becomes

$$\max_{\hat{x}_{t-L|t} \in \mathcal{X}, \hat{d}_{t-L:t-1|t} \in \mathcal{D}} J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{d}_{t-L:t-1|t}, y_{t-L:t}),$$

where now the optimization criterion only contains a finite number of terms that recede as the current time t advances:

$$\begin{aligned} J_t(x_{t-L}, u_{t-L:t-1}, d_{t-L:t-1}, y_{t-L:t}) &\coloneqq -\sum_{s=t-L}^{t-1} \rho_s(d_s) \\ - \sum_{s=t-L}^{t} \eta_s\big(y_s - g\varphi(s; t-L, x_{t-L}, u_{t-L:s-1}, d_{t-L:s-1})\big), \end{aligned}$$

which is essentially the moving horizon estimation problem considered in [14, 15].

III. Main Results

We now show that both for the infinite-horizon and finite-horizon cases introduced in Sections II-A and II-B, respectively, the control law (7) leads to boundedness of the state of the closed-loop system under appropriate assumptions, which we discuss next.

A necessary condition for the implementation of the control law (7) is that the outer minimizations in (6) or (9) lead to finite values for the optima that are achieved at specific sequences $\hat{u}^*_{t:\infty|t} \in \mathcal{U}$, $t \in \mathbb{Z}_{\geqslant 0}$. However, for the stability results in this section we actually ask for the existence of a saddle-point solution to the min-max optimizations in (6) or (9), which is a common requirement in game theoretical approaches to control design [24]:

Assumption 1 (Saddle-point): The min-max optimization (9) always has a saddle-point solution for which the min and max commute. Specifically, for every time $t \in \mathbb{Z}_{\geqslant 0}$, past control input sequence $u_{t-L:t-1} \in \mathcal{U}$, and past measured output sequence $y_{t-L:t} \in \mathcal{Y}$, there exists a finite scalar $J^*_t(u_{t-L:t-1}, y_{t-L:t}) \in \mathbb{R}$, an initial condition $\hat{x}^*_{t-L|t} \in \mathcal{X}$, and sequences $\hat{u}^*_{t:t+T-1|t} \in \mathcal{U}$, $\hat{d}^*_{t-L:t+T-1|t} \in \mathcal{D}$ such that

$$
\begin{aligned}
&J^*_t(u_{t-L:t-1}, y_{t-L:t}) \\
&= J_t(\hat{x}^*_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, \hat{d}^*_{t-L:t+T-1|t}, y_{t-L:t}) \\
&= \min_{\hat{u}_{t:t+T-1|t}\in\mathcal{U}} \max_{\hat{x}_{t-L|t}\in\mathcal{X}, \hat{d}_{t-L:t+T-1|t}\in\mathcal{D}} \\
&J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}) \\
&= \max_{\hat{x}_{t-L|t}\in\mathcal{X}, \hat{d}_{t-L:t+T-1|t}\in\mathcal{D}} \\
&J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}) \qquad (11a) \\
&= \max_{\hat{x}_{t-L|t}\in\mathcal{X}, \hat{d}_{t-L:t+T-1|t}\in\mathcal{D}} \min_{\hat{u}_{t:t+T-1|t}\in\mathcal{U}} \\
&J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}) \\
&= \min_{\hat{u}_{t:t+T-1|t}\in\mathcal{U}} \\
&J_t(\hat{x}^*_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}^*_{t-L:t+T-1|t}, y_{t-L:t}) \qquad (11b) \\
&< \infty. \qquad (11c)
\end{aligned}
$$

For the infinite-horizon case (6), the integer T in (11a)–(11b) should be replaced by ∞, and the integer $t-L$ should be replaced by 0. □

Assumption 1 presumes an appropriate form of *observability/detectability* adapted to the criterion $\sum_{s=t}^{t+T-1} c_s(x_s, u_s, d_s)$ because (11a) implies that, for every initial condition $\hat{x}_{t-L|t} \in \mathcal{X}$ and disturbance sequence $\hat{d}_{t-L:t+T-1|t} \in \mathcal{D}$,

$$
\begin{aligned}
&c\varphi(t; t-L, \hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t|t}, \hat{d}_{t-L:t|t}) \\
&\qquad \leqslant J^*_t(u_{0:t-1}, y_{0:t}) + \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}_{s|t}) \\
&\qquad + \sum_{s=t-L}^{t} \eta_s\big(y_s - g\varphi(s; t-L, \hat{x}_{t-L|t}, u_{t-L:s-1}, \hat{d}_{t-L:s-1|t})\big). \qquad (12)
\end{aligned}
$$

This means that we can bound the size of the *current* state using past outputs and past/future input disturbances. For the infinite-horizon case, Assumption 1 also presumes an appropriate form of *controllability/stabilizability* adapted to the criterion $\sum_{s=t}^{\infty} c_s(x_s, u_s, d_s)$ because (11a) implies that the *future* control sequence $\hat{u}^*_{t:\infty|t} \in \mathcal{U}$ is able to keep "small" the size of *future* states as long as the noise and disturbance remain "small". For the finite-horizon case, a subsequent assumption is needed to ensure controllability.

A. Infinite-Horizon Online Optimization

The following theorem is the main result of this section and provides a bound that can be used to prove boundedness of the state when the control signal is constructed using the *infinite-horizon* criterion (5).

Theorem 1 (Infinite-horizon cost-to-go bound): Suppose that Assumption 1 holds. Then, for every $t \in \mathbb{Z}_{\geqslant 0}$, the trajectories of the process (1) with control (7) defined by the infinite-horizon optimization (6) satisfy

$$
c\varphi(t; 0, x_0, u_{0:t}, d_{0:t}) \leqslant J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \eta_s(n_s) + \sum_{s=0}^{t} \rho_s(d_s). \qquad (13)
$$

□

Next we discuss the implications of Theorem 1 in terms of establishing bounds on the state of the closed-loop system, asymptotic stability, and the ability of the closed-loop to asymptotically track desired trajectories.

1) State boundedness and asymptotic stability: When we select criterion (5), for which there exists a class $\mathcal{K}_\infty$ function $\alpha(\cdot)$ and class $\mathcal{K}$ functions[2] $\beta(\cdot), \delta(\cdot)$ such that

$$
\begin{aligned}
c_t(x, u, d) \geqslant \alpha(\|x\|), \quad \eta_t(n) \leqslant \beta(\|n\|), \quad \rho_t(d) \leqslant \delta(\|d\|), \\
\forall x \in \mathbb{R}^{n_x}, u \in \mathbb{R}^{n_u}, d \in \mathbb{R}^{n_d}, n \in \mathbb{R}^{n_n},
\end{aligned}
$$

we conclude from (13) that, along trajectories of the closed-loop system, the following inequality holds for all $t \in \mathbb{Z}_{\geqslant 0}$:

$$
\alpha(\|x_t\|) \leqslant J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \beta(\|n_s\|) + \sum_{s=0}^{t} \delta(\|d_s\|). \qquad (14)
$$

This provides a bound on the state provided that the noise and disturbances are "vanishing," in the sense that

$$
\sum_{s=0}^{\infty} \beta(\|n_s\|) < \infty, \qquad \sum_{s=0}^{\infty} \delta(\|d_s\|) < \infty.
$$

Theorem 1 also provides bounds on the state for non-vanishing noise and disturbances when we use exponentially time-weighted functions $c_t(\cdot)$, $\eta_t(\cdot)$, and $\rho_t(\cdot)$ that satisfy

$$
c_t(x, u, d) \geqslant \lambda^{-t} \alpha(\|x\|),
$$

[2]A function $\alpha : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ is said to belong to class $\mathcal{K}$ if it is continuous, zero at zero, and strictly increasing and is said to belong to class $\mathcal{K}_\infty$ if it belongs to class $\mathcal{K}$ and is unbounded.

$$\eta_t(n) \leqslant \lambda^{-t}\beta(\|n\|),$$
$$\rho_t(d) \leqslant \lambda^{-t}\delta(\|d\|), \tag{15}$$

for all $x \in \mathbb{R}^{n_x}, u \in \mathbb{R}^{n_u}, d \in \mathbb{R}^{n_d}, n \in \mathbb{R}^{n_n}$ and some $\lambda \in (0,1)$. In this case we conclude from (13) that for all $t \in \mathbb{Z}_{\geqslant 0}$,

$$\alpha(\|x_t\|) \leqslant \lambda^t J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \lambda^{t-s}\beta(\|n_s\|) + \sum_{s=0}^{t} \lambda^{t-s}\delta(\|d_s\|).$$

Therefore, x_t remains bounded provided that the measurement noise n_t and the unmeasured disturbance d_t are both uniformly bounded. Moreover, $\|x_t\|$ converges to zero as $t \to \infty$, when the noise and disturbances vanish asymptotically. We have proved the following:

Corollary 1: Suppose that Assumption 1 holds and also that (15) holds for a class $\mathcal{K}_\infty$ function $\alpha(\cdot)$, class $\mathcal{K}$ functions $\beta(\cdot), \delta(\cdot)$, and $\lambda \in (0,1)$. Then, for every initial condition x_0, uniformly bounded measurement noise sequence $n_{0:\infty}$, and uniformly bounded disturbance sequence $d_{0:\infty}$, the state x_t remains uniformly bounded along the trajectories of the process (1) with control (7) defined by the infinite-horizon optimization (6). Moreover, when d_t and n_t converge to zero as $t \to \infty$, the state x_t also converges to zero. □

Remark 3 (Time-weighted criteria): The exponentially time-weighted functions (15) typically arise from a criterion of the form

$$\sum_{s=t}^{\infty} \lambda^{-s} c(x_s, u_s, d_s) - \sum_{s=0}^{t} \lambda^{-s}\eta(n_s) - \sum_{s=0}^{\infty} \lambda^{-s}\rho(d_s)$$

that weight the future more than the past. In this case, (15) holds for functions $\alpha(\cdot)$, $\beta(\cdot)$, and $\delta(\cdot)$ such that $c(x,u,d) \geqslant \alpha(\|x\|)$, $\eta(n) \leqslant \beta(\|n\|)$, and $\rho(d) \leqslant \delta(\|d\|)$, $\forall x,u,d,n$. □

2) Reference tracking: When the control objective is for the state x_t to follow a given trajectory z_t, the optimization criterion can be selected of the form

$$\sum_{s=t}^{\infty} \lambda^{-s} c(x_s - z_s, u_s, d_s) - \sum_{s=0}^{t} \lambda^{-s}\eta(n_s) - \sum_{s=0}^{\infty} \lambda^{-s}\rho(d_s),$$

with $c(x,u,d) \geqslant \alpha(\|x\|)$, $\forall x,u,d$ for some class $\mathcal{K}_\infty$ function α and $\lambda \in (0,1)$. In this case, we conclude from (13) that, for all $t \in \mathbb{Z}_{\geqslant 0}$,

$$\alpha(\|x_t - z_t\|) \leqslant \lambda^t J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \lambda^{t-s}\eta(n_s) + \sum_{s=0}^{t} \lambda^{t-s}\rho(d_s),$$

which allows us to conclude that x_t converges to z_t as $t \to \infty$ when both n_t and d_t are vanishing sequences, and also that, when these sequences are "ultimately small", the tracking error $x_t - z_t$ will converge to a small value.

B. Finite-Horizon Online Optimization

To establish state boundedness under the control (7) defined by the *finite-horizon* optimization criterion (10), one needs additional assumptions regarding the dynamics and the terminal cost $q_t(\cdot)$.

Assumption 2 (Reversible Dynamics): For every $t \in \mathbb{Z}_{\geqslant 0}$, $x_{t+1} \in \mathcal{X}$, and $u_t \in \mathcal{U}$, there exists a state $\tilde{x}_t \in \mathcal{X}$ and a disturbance $\tilde{d}_t \in \mathcal{D}$ such that

$$x_{t+1} = f_t(\tilde{x}_t, u_t, \tilde{d}_t). \tag{16}$$

□

Assumption 3 (ISS-control Lyapunov function): The terminal cost $q_t(\cdot)$ is an *ISS-control Lyapunov function*, in the sense that, for every $t \in \mathbb{Z}_{\geqslant 0}$, $x \in \mathcal{X}$, there exists a control $u \in \mathcal{U}$ such that for all $d \in \mathcal{D}$,

$$q_{t+1}\big(f_t(x,u,d)\big) - q_t(x) \leqslant -c_t(x,u,d) + \rho_t(d). \tag{17}$$

□

Assumption 2 is very mild and essentially means that the sets of disturbances $\mathcal{D}$ and past states $\mathcal{X}$ are sufficiently rich to allow for a jump to any future state in $\mathcal{X}$.

Assumption 3 plays the role of a common assumption in model predictive control, namely that the terminal cost must be a control Lyapunov function for the closed-loop [2]. In the absence of the disturbance d_t, (17) would mean that $q_t(\cdot)$ could be viewed as a control Lyapunov function that decreases along system trajectories for an appropriate control input u_t [28]. With disturbances, $q_t(\cdot)$ needs to be viewed as an ISS-control Lyapunov function that satisfies an ISS stability condition for the disturbance input d_t and an appropriate control input u_t [29].

Remark 4: When the dynamics are linear, for instance, Assumption 2 is satisfied if the state-space A matrix has no eigenvalues at the origin (e.g., if it results from the time-discretization of a continous-time system). When, the dynamics are linear and the cost function is quadratic, a terminal cost $q_t(\cdot)$ satisfying Assumption 3 is typically found by solving a system of linear matrix inequalities. □

We are now ready to state the finite-horizon counter-part to Theorem 1.

Theorem 2 (Finite-horizon cost-to-go bound): Suppose that Assumptions 1, 2, and 3 hold. Then, for every $t \in \mathbb{Z}_{\geqslant 0}$, there exist vectors $\tilde{d}_s \in \mathbb{R}^{n_d}$, $\tilde{n}_s \in \mathbb{R}^{n_n}$, for which

$$\eta_s(\tilde{n}_s) < \infty, \quad \rho_s(\tilde{d}_s) < \infty, \quad \forall s \in \{0, 1, ..., t-L-1\},$$

and the trajectories of the process (1) with control (7) defined by the finite-horizon optimization (9) satisfy

$$c\varphi(t; t-L, x_{t-L}, u_{t-L:t}, d_{t-L:t}) \leqslant J_0^*(y_0) + \sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s) + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s) + \sum_{s=t-L}^{t} \eta_s(n_s) + \sum_{s=t-L}^{t} \rho_s(d_s). \tag{18}$$

□

The terms $\sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s) + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s)$ in the right-hand side of (18) can be thought of as the *arrival cost* that appears in the MHE literature to capture the quality of the estimate at the beginning of the current estimation window [14]. For the *extended-value extensions* in (3), these terms can be bounded by

$$\sum_{s=0}^{t-L-1} \sup_{n_s \in \mathcal{N}} \bar{\eta}_s(n_s) + \sum_{s=0}^{t-L-1} \sup_{d_s \in \mathcal{D}} \bar{\rho}_s(d_s).$$

Since (13) and (18) provide nearly identical bounds, the discussion presented after Theorem 1 regarding state boundedness and reference tracking applies also to the finite-horizon case, so we do not repeat it here. Proofs of these results can be found in the technical report [30].

IV. Validation through Simulation

To implement the control law (7) we need to find the control sequence $\hat{u}^*_{t:t+T-1|t} \in \mathcal{U}$ that achieves the outer minimizations in (9). In view of Assumption 1, the desired control sequence must be part of the saddle-point defined by (11a)–(11b). It turns out that, from the perspective of numerically computing this saddle-point, it is more convenient to use the following equivalent characterization of the saddle-point:

$$
\begin{aligned}
&-J^*_t(u_{t-L:t-1}, y_{t-L:t}) = \min_{\hat{x}_{t-L|t}\in\mathcal{X}, \hat{d}_{t-L:t+T-1|t}\in\mathcal{D}} \\
&-J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}),
\end{aligned} \tag{19a}
$$

$$
\begin{aligned}
&J^*_t(u_{t-L:t-1}, y_{t-L:t}) = \min_{\hat{u}_{t:t+T-1|t}\in\mathcal{U}} \\
&J_t(\hat{x}^*_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}^*_{t-L:t+T-1|t}, y_{t-L:t})
\end{aligned} \tag{19b}
$$

(see, e.g., [24]). We introduce the "$-$" sign in (19a) simply to obtain two minimizations, instead of a maximization and one minimization, which somewhat simplifies the presentation.

Since the process dynamics (1) has a unique solution for any given initial condition, control input, and unmeasured disturbance, the coupled optimizations in (19) can be rewritten as

$$
\begin{aligned}
&-J^*_t(u_{t-L:t-1}, y_{t-L:t}) = \\
&\min_{(\hat{d}_{t-L:t+T-1|t}, \bar{x}_{t-L:t+T|t})\in\bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}]} \\
&\quad -\sum_{s=t}^{t+T-1} c_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s) - q_{t+T}(\bar{x}_{t+T}) \\
&\quad + \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\bar{x}_s)\big) + \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}_s),
\end{aligned} \tag{20a}
$$

$$
\begin{aligned}
&J^*_t(u_{t-L:t-1}, y_{t-L:t}) = \\
&\min_{(\hat{u}_{t:t+T-1|t}, \tilde{x}_{t-L+1:t+T|t})\in\bar{\mathcal{U}}[\bar{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}]} \\
&\quad \sum_{s=t}^{t+T-1} c_s(\tilde{x}_s, \hat{u}_s, \hat{d}^*_s) + q_{t+T}(\tilde{x}_{t+T}) \\
&\quad - \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\tilde{x}_s)\big) - \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}^*_s),
\end{aligned} \tag{20b}
$$

where

$$
\begin{aligned}
\bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}] \coloneqq \Big\{&(\hat{d}_{t-L:t+T-1|t}, \bar{x}_{t-L:t+T|t}) : \\
&\hat{d}_{t-L:t+T-1|t} \in \mathcal{D},\ \bar{x}_{t-L:t+T|t} \in \mathcal{X}, \\
&\bar{x}_{s+1} = f_s(\bar{x}_s, u_s, \hat{d}_s), \forall s \in \{t-L, ..., t-1\}, \\
&\bar{x}_{s+1} = f_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s), \forall s \in \{t, ..., t+T-1\}\Big\},
\end{aligned} \tag{21a}
$$

$$
\begin{aligned}
\bar{\mathcal{U}}[\bar{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}] \coloneqq \Big\{&(\hat{u}_{t:t+T-1|t}, \tilde{x}_{t-L+1:t+T|t}) : \\
&\hat{u}_{t:t+T-1|t} \in \mathcal{U},\ \tilde{x}_{t-L+1:t+T|t} \in \mathcal{X}, \\
&\tilde{x}_{t-L+1} = f_s(\bar{x}^*_{t-L}, u_{t-L}, \hat{d}^*_{t-L}), \\
&\tilde{x}_{s+1} = f_s(\tilde{x}_s, u_s, \hat{d}^*_s), \forall s \in \{t-L+1, ..., t-1\}, \\
&\tilde{x}_{s+1} = f_s(\tilde{x}_s, \hat{u}_s, \hat{d}^*_s), \forall s \in \{t, ..., t+T-1\}\Big\}.
\end{aligned} \tag{21b}
$$

Essentially, in each of the optimizations in (19) we introduced the values of the state as additional optimization variables that are constrained by the system dynamics. While this introduces additional optimization variables, it avoids the need to explicitly evaluate the solution $\varphi(t; t-L, x_{t-L}, u_{t-L:t-1}, d_{t-L:t-1})$ that appears in the original optimizations (19) and that can be numerically poorly conditioned, e.g., for systems with unstable dynamics.

While solving either (19) or (20) gives the same results, we choose to solve the latter because it generally leads to simpler optimization problems. An efficient primal-dual-like interior-point method was developed in order to numerically compute solutions to this min-max optimization. A description of this algorithm is of interest on its own and can be found in [30]. The following example uses this approach.

Example 1 (Flexible beam): Consider a single-link flexible beam like the one described in [31], where the control objective is to regulate the mass on the tip of the beam to a desired reference trajectory. The control input is the applied torque at the base, and the outputs are the tip's position, the angle at the base, the angular velocity of the base, and a strain gauge measurement collected around the middle of the beam, respectively.

An approximate linearized discrete-time state-space model of the dynamics, with a sampling time $T_s \coloneqq 1$ second, is given by $x_{t+1} = Ax_t + B(u_t + d_t)$, $y_t = Cx_t + n_t$, where d_t is a disturbance, n_t is measurement noise, and the system matrices are given by

$$
A = \begin{bmatrix}
1.0 & 1.016 & -0.676 & -1.084 & 1.0 & 0.585 & 0.233 & 0.032 \\
0 & -0.665 & 1.241 & 1.783 & 0 & 0.042 & -0.288 & -0.023 \\
0 & 0.009 & -0.439 & 0.143 & 0 & -0.002 & -0.012 & 0.007 \\
0 & 0.001 & 0.014 & 0.308 & 0 & -0.000 & 0.001 & 0.001 \\
0 & 1.264 & -37.070 & 10.581 & 1.0 & 1.016 & -0.676 & -1.084 \\
0 & -2.109 & 59.920 & -16.883 & 0 & -0.665 & 1.241 & 1.783 \\
0 & 0.413 & 9.156 & -3.695 & 0 & 0.009 & -0.439 & 0.143 \\
0 & -0.012 & -0.371 & -3.929 & 0 & 0.001 & 0.014 & 0.308
\end{bmatrix},
$$

$$
B = \begin{bmatrix} 0.800 & -0.797 & 0.003 & 0.001 & 1.327 & -1.163 & 0.197 & -0.006 \end{bmatrix}',
$$

$$
C = \begin{bmatrix}
1.13 & 0.7225 & -0.2028 & 0.1220 & 0 & 0 & 0 & 0 \\
1.0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1.0 & 0 & 0 & 0 \\
0 & 0.9282 & -12.001 & -35.294 & 0 & 0 & 0 & 0
\end{bmatrix}. \tag{22}
$$

This matrix A has a double eigenvalue at 1 with a single independent eigenvector. Therefore this is an unstable system.

The optimal control input is found by solving the following optimization problem

$$
\begin{aligned}
&\min_{u_{t:t+T-1|t}\in\mathcal{U}} \ \max_{x_{t-L}\in\mathcal{X}, d_{t-L:t+T-1|t}\in\mathcal{D}} \|x_{t:t+T} - r_{t:t+T}\|^2 \\
&\quad + \lambda_u \|u_{t:t+T-1}\|^2 - \lambda_d \|d_{t-L:t+T-1}\|^2 - \lambda_n \|n_{t-L:t}\|^2,
\end{aligned}
$$

where $\|\cdot\|$ is the Euclidean norm, $\mathcal{U} \coloneqq \{u_t \in \mathbb{R} | -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} \coloneqq \mathbb{R}^8$, and $\mathcal{D} \coloneqq \{d_t \in \mathbb{R} | -d_{max} \leqslant d_t \leqslant d_{max}\}$. The numerical computation of solutions to the min-max optimization was performed using the algorithm described in [30].

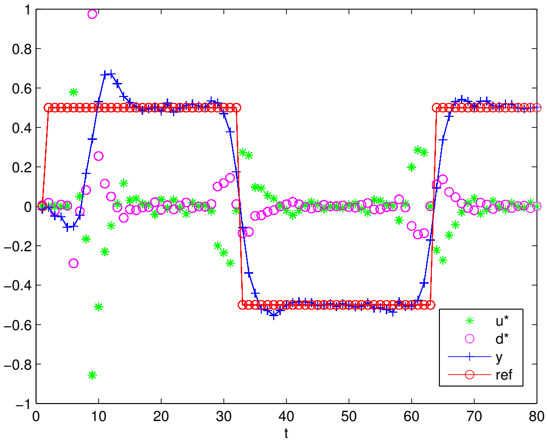

Fig. 1. Simulation results of Flexible Beam example. The reference is in red, the measured output in blue, the control sequence in green, and the disturbance sequence in magenta.

The results depicted in Figure 1 show the response of the closed loop system under the control law (7) when our goal is to regulate the mass at the tip of the beam to a desired reference $r(t) \coloneqq \alpha \operatorname{sgn}(\sin(\omega t))$ with $\alpha = 0.5$ and $\omega = 0.1$. The other parameters in the optimization have values $\lambda_u = 1$, $\lambda_d = 2$, $\lambda_n = 100$, $L = 5$, $T = 5$, $u_{max} = 1$, $d_{max} = 1$. The state of the system starts close to zero and evolves with zero control input and small random disturbance input until time $t = 6$, at which time the optimal control input (7) started to be applied along with the optimal worst-case disturbance $d^*_{t|t}$ obtained from the min-max optimization. The noise process n_t was selected to be a zero-mean Gaussian independent and identically distributed random process with standard deviation of 0.01.

V. Conclusions

We presented an output-feedback approach to nonlinear model predictive control using moving horizon state estimation. Solutions to the combined control and state estimation problems were found by solving a single min-max optimization problem. Under the assumption that a saddle-point solution exists (which presumes appropriate forms of observability and controllability), Theorem 1 gives bounds on the state of the system and the tracking error for reference tracking problems. Similar results are given in Theorem 2 for the finite-horizon case under the additional assumptions of reversible dynamics and a terminal cost that is an ISS-control Lyapunov function with respect to the disturbance input.

Directions for future work include investigating under what conditions, such as controllability and observability, a saddle-point solution exists. Results specific to certain types of uncertainty, noise, and disturbances, such as model uncertainty, may also be investigated.

References

[1] M. Morari and J. H Lee, "Model predictive control: past, present and future," *Computers & Chemical Engineering*, vol. 23, no. 4, pp. 667–682, 1999.

[2] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.

[3] J. B. Rawlings and D. Q. Mayne, *Model Predictive Control: Theory and Design*. Nob Hill Publishing, 2009.

[4] P. J. Campo and M. Morari, "Robust model predictive control," in *American Control Conference, 1987*, pp. 1021–1026, IEEE, 1987.

[5] J. a. Lee and Z. Yu, "Worst-case formulations of model predictive control for systems with bounded parameters," *Automatica*, vol. 33, no. 5, pp. 763–781, 1997.

[6] A. Bemporad and M. Morari, "Robust model predictive control: A survey," in *Robustness in identification and control*, pp. 207–226, Springer, 1999.

[7] L. Magni, G. De Nicolao, R. Scattolini, and F. Allgöwer, "Robust model predictive control for nonlinear discrete-time systems," *International Journal of Robust and Nonlinear Control*, vol. 13, no. 3-4, pp. 229–246, 2003.

[8] P. Scokaert and D. Mayne, "Min-max feedback model predictive control for constrained linear systems," *Automatic Control, IEEE Transactions on*, vol. 43, no. 8, pp. 1136–1142, 1998.

[9] A. Bemporad, F. Borrelli, and M. Morari, "Min-max control of constrained uncertain discrete-time linear systems," *Automatic Control, IEEE Transactions on*, vol. 48, no. 9, pp. 1600–1606, 2003.

[10] H. Chen, C. Scherer, and F. Allgöwer, "A game theoretic approach to nonlinear robust receding horizon control of constrained systems," in *American Control Conference, 1997. Proceedings of the 1997*, vol. 5, pp. 3073–3077, IEEE, 1997.

[11] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel, "Nominally robust model predictive control with state constraints," *Automatic Control, IEEE Transactions on*, vol. 52, no. 10, pp. 1856–1870, 2007.

[12] J. B. Rawlings and B. R. Bakshi, "Particle filtering and moving horizon estimation," *Computers & chemical engineering*, vol. 30, no. 10, pp. 1529–1541, 2006.

[13] F. Allgöwer, T. A. Badgwell, J. S. Qin, J. B. Rawlings, and S. J. Wright, "Nonlinear predictive control and moving horizon estimation-an introductory overview," in *Advances in control*, pp. 391–449, Springer, 1999.

[14] C. V. Rao, J. B. Rawlings, and D. Q. Mayne, "Constrained state estimation for nonlinear discrete-time systems: Stability and moving horizon approximations," *Automatic Control, IEEE Transactions on*, vol. 48, no. 2, pp. 246–258, 2003.

[15] A. Alessandri, M. Baglietto, and G. Battistelli, "Moving-horizon state estimation for nonlinear discrete-time systems: New stability results and approximation schemes," *Automatica*, vol. 44, no. 7, pp. 1753–1765, 2008.

[16] J. Löfberg, "Towards joint state estimation and control in minimax MPC," in *Proceedings of 15th IFAC World Congress, Barcelona, Spain*, 2002.

[17] D. Q. Mayne, S. Raković, R. Findeisen, and F. Allgöwer, "Robust output feedback model predictive control of constrained linear systems: Time varying case," *Automatica*, vol. 45, no. 9, pp. 2082–2087, 2009.

[18] D. Sui, L. Feng, and M. Hovd, "Robust output feedback model predictive control for linear systems via moving horizon estimation," in *American Control Conference, 2008*, pp. 453–458, IEEE, 2008.

[19] R. Findeisen, L. Imsland, F. Allgöwer, and B. A. Foss, "State and output feedback nonlinear model predictive control: An overview," *European Journal of Control*, vol. 9, no. 2, pp. 190–206, 2003.

[20] L. Imsland, R. Findeisen, E. Bullinger, F. Allgöwer, and B. A. Foss, "A note on stability, robustness and performance of output feedback nonlinear model predictive control," *Journal of Process Control*, vol. 13, no. 7, pp. 633–644, 2003.

[21] M. Lazar, D. Muñoz de la Peña, W. Heemels, and T. Alamo, "On input-to-state stability of min–max nonlinear model predictive control," *Systems & Control Letters*, vol. 57, no. 1, pp. 39–48, 2008.

[22] D. Limon, T. Alamo, D. Raimondo, D. M. de la Pena, J. Bravo, A. Ferramosca, and E. Camacho, "Input-to-state stability: a unifying framework for robust model predictive control," in *Nonlinear model predictive control*, pp. 1–26, Springer, 2009.

[23] D. M. Raimondo, D. Limon, M. Lazar, L. Magni, and E. Camacho, "Min-max model predictive control of nonlinear systems: A unifying overview on stability," *European Journal of Control*, vol. 15, no. 1, pp. 5–21, 2009.

[24] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. London: Academic Press, 1995.

[25] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

[26] L. Magni and R. Scattolini, "Robustness and robust design of MPC for nonlinear discrete-time systems," in *Assessment and future directions of nonlinear model predictive control*, pp. 239–254, Springer, 2007.

[27] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, pp. 3851–3861, IEEE, 2012.

[28] E. D. Sontag, "Control-lyapunov functions," in *Open problems in mathematical systems and control theory*, pp. 211–216, Springer, 1999.

[29] D. Liberzon, E. D. Sontag, and Y. Wang, "Universal construction of feedback laws achieving ISS and integral-ISS disturbance attenuation," *Systems & Control Letters*, vol. 46, no. 2, pp. 111–127, 2002.

[30] D. A. Copp and J. P. Hespanha, "Nonlinear output-feedback model predictive control with moving horizon estimation," tech. rep., Univ. California, Santa Barbara, 2014. `http://www.ece.ucsb.edu/~hespanha/techrep.html`.

[31] E. Schmitz, *Experiments on the End-Point Position Control of a Very Flexible One-Link Manipulator*. PhD thesis, Stanford University, 1985.