Nonlinear Output-Feedback Model Predictive Control with Moving Horizon Estimation (Technical Report)

David A. Copp and João P. Hespanha*

December 9, 2014 †

Abstract

We introduce an output-feedback approach to model predictive control that combines state estimation and control into a single min-max optimization. Like in the more common state-feedback MPC, this approach allows one to incorporate explicit constraints on the control input and state. In addition, it allows one to incorporate any known constraints on disturbances and noise. 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 observability of the nonlinear system and the use of a terminal cost that is an ISS-control Lyapunov function with respect to a disturbance input. We also introduce a primal-dual-like interior-point method to solve the min-max optimization that arises in our approach. Under appropriate convexity assumptions, this method is guaranteed to terminate at a global solution. However, simulation results show that it also converges rapidly in many problems that are severely nonconvex. This report includes a few representative examples that demonstrate the applicability of the approach in systems that are high-dimensional, nonlinear in their dynamics and/or measurements, and that have significant dynamic uncertainty.

1 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 [1–4]. In MPC, the current state of the plant to be controlled is used as an initial condition from which an online optimization is solved. This optimization yields an optimal control sequence from which the first control action in the sequence is selected and applied to the plant. Then, at each sampling time this technique is repeated. A nice tutorial overview of MPC is given in [5].

MPC is often attractive in many applications because it can explicitly handle hard state and input constraints, but a downside to MPC is the computational complexity involved in solving these problems rapidly online. Because MPC problems require the solution of an optimization problem at each sampling time, efficient numerical methods for solving these problems are imperative for effective control. In the past,

*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

†This report was updated on September 8, 2016 to correct for inaccuracies and typos.

MPC has been popular in many industries where plant dynamics are slow enough to accommodate the time necessary to numerically compute solutions online. Now that computational efficiency has increased, MPC is penetrating even more areas in industry. For a survey of MPC applications in industry, see [6].

Besides computational complexity, there are other things to consider when using MPC such as robustness to model uncertainty, input disturbances, and measurement noise. The study of these topics are known as robust, worst-case, or min-max MPC. Initial results on these topics are discussed in works such as [7–10]. In order to alleviate problems from uncertainties, noise, and disturbances, 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 robust output-feedback MPC in which the use of an independent algorithm for state estimation is required. Examples of algorithms for state estimation include observers, filters, and moving horizon estimation, some of which are discussed in [11]. Of these methods, moving horizon estimation (MHE) is attractive for use with MPC because it 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 are both discussed in [12]. A useful overview of constrained nonlinear moving horizon state estimation is given in [13].

Thus far, results on the stability of output-feedback control schemes based on MPC and MHE (especially for nonlinear systems) are limited. 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 at each sampling time. This assumption is a common requirement in game theoretical approaches to control design [14] and presumes appropriate forms of observability and controllability of the closed-loop system. For the finite-horizon case, we require an additional observability assumption and that there exists a terminal cost that is an ISS-control Lyapunov function with respect to a disturbance input.

Several algorithms are available to numerically solve the class of min-max optimization problems that we discuss here. A few methods are discussed in [15] and [16] and include sequential quadratic programming, interior-point methods, and others. We propose a primal-dual-like interior-point algorithm to solve this min-max optimization. In the case of solving convex problems, our algorithm is guaranteed to find the global solution if it converges due to satisfying first-order optimality conditions. In the nonconvex case, we cannot guarantee that the result is a global minimum/maximum, but we show in simulations of non-convex examples that a local minimum/maximum may still be found. We present several constrained linear and nonlinear examples of output-feedback MPC with MHE, and use our primal-dual-like interior-point algorithm to solve them. These examples show robust trajectory tracking in the presence of additive disturbances and measurement noise and effective adaptive control in the presence of uncertain model parameters. We do not investigate the numerical performance of our method but show that it converges to the correct solution and that it is reliable even when solving nonconvex problems.

The paper is organized as follows. First, we briefly describe related work that has been done in the areas of model predictive control, moving horizon estimation, numerical methods for min-max optimization problems, and specifically primal-dual methods. In Section 2, we formulate the control problem we would like

to solve and discuss its relationship to MPC and MHE. In Section 3, we state the main closed-loop stability results. Simulation results are presented in Section 6, and we provide some conclusions and directions for future research in Section 7.

Related Work

Model predictive control, moving horizon estimation, and numerical optimization are each large areas of study, so now we mention some related work to narrow our focus. As discussed in the introduction, the study of model predictive control is quite mature as evidenced by [1–4]. Robust and worst-case MPC is initially discussed in works such as [7–10]. Min-max MPC for constrained linear systems is considered in [17] and [18], and a game theoretic approach for robust constrained nonlinear MPC is proposed in [19]. Recent studies of input-to-state stability of min-max MPC can be found in [20–22], however these references do not investigate the use of output-feedback. Nominal or inherent robustness of MPC has also been studied in [3, 23].

Because MPC and MHE problems can be formulated as similar optimization problems, and because output-feedback MPC requires some form of state estimation, during the same time that many important results on MPC were developed, parallel work was being done on MHE. Nice overviews of constrained linear and nonlinear moving horizon state estimation can be found in [13, 24, 25]. Recent results regarding stability of MHE can be found in [26]. Nonlinear MPC and MHE are both discussed in [12]. Some joint stability results for state estimation and control are given in [27], but again, output-feedback MPC is not considered. Recently, more work has been done on output-feedback MPC. A survey including some nonlinear results is given in [28]. Results on robust output-feedback MPC for constrained linear systems can be found in [29] using a state observer for estimation, and in [30] using MHE for estimation. Fewer results are available for nonlinear output-feedback MPC, although notable exceptions are [3, 31]. Recent studies of input-to-state stability of min-max MPC can be found in [20–22], however, these references do not investigate the use of output-feedback.

Numerical optimization is an extensive field involving the derivation of methods to numerically solve optimization problems, such as those that appear in MPC and MHE, efficiently and reliably. A great place to start studying convex optimization problems and methods to solve them (including interior-point and primal-dual interior-point methods) is in the book [32]. Work regarding interior-point methods can be found in [33, 34] and primal-dual interior-point methods in particular in [35]. The application of interior-point algorithms as a method to solve MPC problems is discussed in [36]. Other early work on efficient numerical methods for solving MPC problems are given in [37], [38], and [35]. Advances in computational efficiency have allowed for the fast solution of MPC problems using online optimization such as in the recent work [39]. The real-time solution of the MHE problem for small dimensional nonlinear models is given in [15]. Considering specifically numerical methods for min-max MPC optimization problems, the authors in [40] set up and solve min-max MPC as a quadratic program. Robust dynamic programming for min-max MPC of constrained uncertain systems is considered in [41], while sequential quadratic programming and interior-point methods for solving nonlinear MPC with MHE problems are discussed in [16]. The particular method that we describe in this paper is inspired by the primal-dual interior-point method for a single optimization given in [42].

2 Problem Formulation

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

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

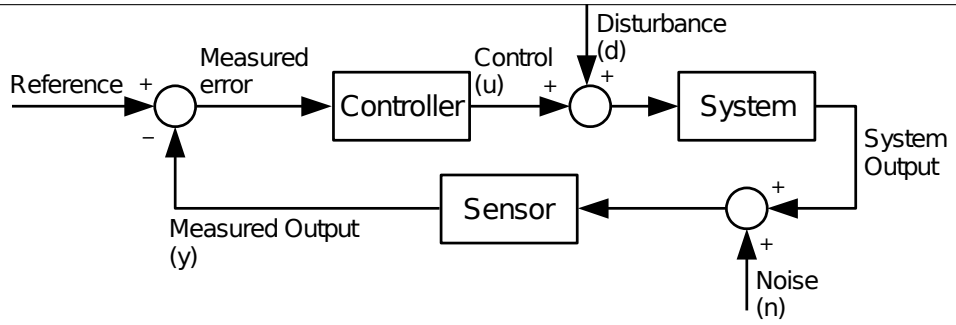

Figure 1: Block diagram depicting the problem formulation.

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 that is known to belong to the set $\mathcal{N} \subset \mathbb{R}^{n_n}$. The signal $y_t \in \mathbb{R}^{n_y}$ denotes the *measured output* that is available for feedback. A block diagram depicting this problem formulation is shown in Figure 1.

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 $n_t \in \mathbb{R}^{n_n}$, $\forall t \in \mathbb{Z}_{\geqslant 0}$. 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* [32] of the form

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

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 (Quadratic case). 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. □

2.1 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 $\{z_{t_0}, z_{t_0+1}, \ldots, z_t\}$. 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

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

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

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

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 the criterion (5) depends, so we optimize this criterion under worst-case assumptions on these variables, leading to the following min-max optimization

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

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

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 variables $\hat{n}_{0:t}$ are not independent optimization variables as they are uniquely determined by the remaining optimization variables and the output equation:

$$\hat{n}_{s|t} = y_s - g_s(\hat{x}_{s|t}), \quad \forall s \in \{0, 1, \dots, t\}.$$

Consequently, the condition $\hat{n}_{0:t|t} \in \mathcal{N}$ can simply be regarded as a constraint on the remaining optimization variables for the (inner) maximization.

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 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:

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

2.2 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

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

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

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

where

$$J_t(x_{t-L}, u_{t-L:t+T-1}, d_{t-L:t+T-1}, y_{t-L:t}) \coloneqq \sum_{s=t}^{t+T-1} c\varphi(s; t-L, x_{t-L}, u_{t-L:s}, d_{t-L:s})$$

$$+ q\varphi(t+T; t-L, x_{t-L}, u_{t-L:t+T-1}, d_{t-L:t+T-1}) - \sum_{s=t-L}^{t} \eta_s\big(y_s - g\varphi(s; t-L, x_{t-L}, u_{t-L:s-1}, d_{t-L:s-1})\big) - \sum_{s=t-L}^{t+T-1} \rho_s(d_s). \quad (10)$$

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).

2.3 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

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

with

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

which is essentially the robust model predictive control with terminal cost considered in [19, 43].

Remark 2 (Economic MPC). It is worth noting that our framework is more general than standard forms of MPC. It can also apply to economic MPC in which the operating cost of the plant is used directly in the MPC objective function [44]. □

2.4 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

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

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

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

which is essentially the moving horizon estimation problem considered in [13, 26].

A depiction of the finite-horizon control and estimation problems is shown in Figure 2.

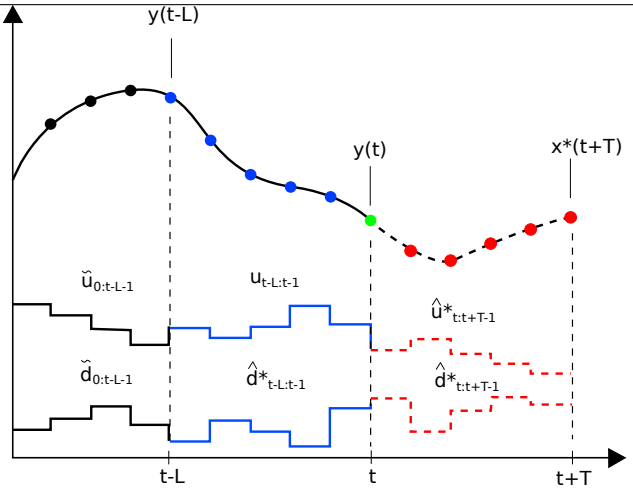

Figure 2: Finite-horizon control and estimation problems. The elements in blue correspond to the MHE problem, and the elements in red correspond to the MPC problem.

3 Main Results

We now show that both for the infinite-horizon and finite-horizon cases introduced in Sections 2.1 and 2.2, 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 assumption 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 [14]:

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

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

$$< \infty. \tag{11c}$$

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 $\hat{d}_{t-L:t+T-1|t} \in \mathcal{D}$,

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

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.

For linear systems and quadratic costs, Assumption 1 is satisfied if the system is observable and the weights in the cost function are chosen appropriately [45].

3.1 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*

$$c\varphi(t; 0, x_0, u_{0:t}, d_{0:t}) \leqslant J^{\infty *}_0(y_0) + \sum_{s=0}^{t} \eta_s(n_s) + \sum_{s=0}^{t} \rho_s(d_s). \tag{13}$$

□

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.

3.1.1 State boundedness and asymptotic stability

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

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

[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 to belong to class $\mathcal{K}_\infty$ if it belongs to class $\mathcal{K}$ and is unbounded.

we conclude from (13) that, along trajectories of the closed-loop system, we have

$$\alpha(\|x_t\|) \leqslant J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \beta(\|n_s\|) + \sum_{s=0}^{t} \delta(\|d_s\|), \quad \forall t \in \mathbb{Z}_{\geqslant 0}. \tag{14}$$

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 can also provide 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

$$c_t(x,u,d) \geqslant \lambda^{-t}\alpha(\|x\|), \quad \eta_t(n) \leqslant \lambda^{-t}\beta(\|n\|), \quad \rho_t(d) \leqslant \lambda^{-t}\delta(\|d\|), \quad \forall x \in \mathbb{R}^{n_x}, u \in \mathbb{R}^{n_u}, d \in \mathbb{R}^{n_d}, n \in \mathbb{R}^{n_n}, \tag{15}$$

for some $\lambda \in (0,1)$, in which case we conclude from (13) that

$$\alpha(\|x_t\|) \leqslant \lambda^t J_0^{\infty *}(y_0) + \sum_{s=0}^{t} \lambda^{t-s}\beta(\|n_s\|) + \sum_{s=0}^{t} \lambda^{t-s}\delta(\|d_s\|), \qquad \forall t \in \mathbb{Z}_{\geqslant 0}.$$

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 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 α, β, and δ such that $c(x,u,d) \geqslant \alpha(\|x\|)$, $\eta(n) \leqslant \beta(\|n\|)$, and $\rho(d) \leqslant \delta(\|d\|)$, $\forall x, u, d, n$. □

3.1.2 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

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

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.

3.2 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 terminal cost $q_t(\cdot)$ as well as observability of the nonlinear dynamics.

Assumption 2 (Observability)**.** There exists a bounded set $\mathcal{N}_{\text{pre}} \subset \mathbb{R}^{n_n}$ such that, for every time $t \in \mathbb{Z}_{\geqslant 0}$, every state $\hat{x}_{t-L:t} \in \mathcal{X}$, and every disturbance and noise sequence, $\hat{d}_{t-L:t} \in \mathcal{D}$ and $\hat{n}_{t-L:t} \in \mathcal{N}$, that are compatible with the applied control input u_s, $s \in \mathbb{Z}_{\geqslant 0}$, and the measured output y_s, $s \in \mathbb{Z}_{\geqslant 0}$, in the sense that

$$\hat{x}_{s+1} = f_s(\hat{x}_s, u_s, \hat{d}_s), \qquad\qquad y_s = g_s(\hat{x}_s) + \hat{n}_s, \tag{16}$$

$\forall s \in \{t-L, t-L+1, \ldots, t\}$, there exists a "predecessor" state estimate $\hat{x}_{t-L-1} \in \mathcal{X}$, disturbance estimate $\hat{d}_{t-L-1} \in \mathcal{D}$, and noise estimate $\hat{n}_{t-L-1} \in \mathcal{N}_{\text{pre}}$ such that (16) also holds for time $s = t-L-1$. □

In essence, Assumption 2 requires the past horizon length L to be sufficiently large so that, by observing the system's inputs and outputs over a past time interval $\{t-L, t-L+1, \ldots, t\}$, one obtains enough information about the initial condition x_{t-L} so that any estimate $\hat{x}_{t-L}$ that is compatible with the observed input/output data is "precise". By "precise," we mean that if one were to observe one additional past input/output pair u_{t-L-1}, y_{t-L-1} just before the original interval, it would be possible to find an estimate $\hat{x}_{t-L-1}$ for the "predecessor" state x_{t-L-1} that would be compatible with the previous estimate $\hat{x}_{t-L}$, that is,

$$\hat{x}_{t-L} = f_{t-L-1}(\hat{x}_{t-L-1}, u_{t-L-1}, \hat{d}_{t-L-1}).$$

This "predecessor" state estimate $\hat{x}_{t-L-1}$ would also be compatible with the measured output at time $t-L-1$ in the sense that the output estimation error lies in the bounded set $\mathcal{N}_{\text{pre}}$:

$$y_{t-L-1} - g_{t-L-1}(\hat{x}_{t-L-1}) \in \mathcal{N}_{\text{pre}}. \tag{17}$$

We do not require the bounded set $\mathcal{N}_{\text{pre}}$ to be the same as the set $\mathcal{N}$ in which the actual noise is known to lie. In fact, the set $\mathcal{N}_{\text{pre}}$ where the "predecessor" output error (17) should lie may have to be made larger than $\mathcal{N}$ to make sure that Assumption 2 holds. For linear systems, it is straightforward to argue that Assumption 2 holds provided that the matrix

$$\begin{bmatrix} C \\ CA \\ \vdots \\ CA^L \end{bmatrix}$$

is full column rank and the set $\mathcal{N}_{\text{pre}}$ is chosen sufficiently large. For nonlinear systems, computing the set $\mathcal{N}_{\text{pre}}$ may be difficult, but fortunately we do not need to compute this set to implement the controller.

Remark 4 (Choosing length of L). Although computing the set $\mathcal{N}_{\text{pre}}$ is not required, how large $\mathcal{N}_{\text{pre}}$ needs to be is essentially determined by the length of the backwards horizon L. As the length of L is increased, equation (16) provides more constraints on the estimates which leads to better estimates and, therefore, a necessarily smaller set $\mathcal{N}_{\text{pre}}$. Therefore, larger L is generally better, but increasing L also increases the computation required to solve (9) as the number of optimization variables increases as well. Thus, a heuristic for choosing L is to make it as large as possible given available computational resources. □

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

$$q_{t+1}\big(f_t(x,u,d)\big) - q_t(x) \leqslant -c_t(x,u,d) + \rho_t(d), \quad \forall d \in \mathcal{D}. \tag{18}$$

□

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 [46]. In the absence of the disturbance d_t, (18) 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 [47]. 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 [48].

Remark 5 (Linear/Quadratic cost). 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. Along any trajectory of the closed-loop system defined by the process* (1) *and the control law* (7) *defined by the finite-horizon optimization* (9)*, we have that*

$$\begin{aligned} &c\varphi(t; t-L, x_{t-L}, u_{t-L:t}, d_{t-L:t}) \\ &\qquad \leqslant J_L^*(u_{0:L-1}, y_{0:L}) + \sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s) + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s) + \sum_{s=t-L}^{t} \eta_s(n_s) + \sum_{s=t-L}^{t} \rho_s(d_s). \end{aligned} \tag{19}$$

for appropriate sequences $\tilde{d}_{0:t-L-1} \in \mathcal{D}$, $\tilde{n}_{0:t-L-1} \in \mathcal{N}_{\text{pre}}$. □

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 (19) 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 [13].

Since (13) and (19) 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.

3.3 Proofs

We now present the proof of Theorem 2. The proof of Theorem 1 is omitted because it is simpler and can be obtained from the proof of Theorem 2 by systematically replacing T by $+\infty$, $t-L$ by 0, and dropping all terms in the arrival cost $\sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s) + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s)$ and terminal cost $q_{(\cdot)}(\cdot)$, wherever they may appear.

Before proving Theorem 2, we introduce a key technical lemma that establishes a monotonicity-like property of the sequence $\{J_t^* : t \in \mathbb{Z}_{\geqslant 0}\}$ computed along solutions to the closed loop.

Lemma 1. *Suppose that Assumptions 1, 2, and 3 hold. Along any trajectory of the closed-loop system defined by the process* (1) *and the control law* (7)*, the sequence* $\{J_t^* : t \in \mathbb{Z}_{\geqslant 0}\}$*, whose existence is guaranteed by Assumption 1, satisfies*

$$J_{t+1}^* - J_t^* \leqslant \eta_{t-L}(\tilde{n}_{t-L}) + \rho_{t-L}(\tilde{d}_{t-L}), \quad \forall t \in \mathbb{Z}_{\geqslant L} \tag{20}$$

for appropriate sequences $\tilde{d}_{0:t-L-1} \in \mathcal{D}$, $\tilde{n}_{0:t-L-1} \in \mathcal{N}_{\text{pre}}$. □

Proof of Lemma 1. From (11b) in Assumption 1 at time $t+1$, we conclude that there exists an initial condition $\hat{x}_{t-L+1|t+1}^* \in \mathcal{X}$ and sequences $\hat{d}_{t-L+1:t+T|t+1}^* \in \mathcal{D}$, $\hat{n}_{t-L+1:t+1|t+1}^* \in \mathcal{N}$ such that

$$J_{t+1}^* = \min_{\hat{u}_{t+1:t+T|t+1} \in \mathcal{U}} J_{t+1}(\hat{x}_{t-L+1|t+1}^*, u_{t-L+1:t}, \hat{u}_{t+1:t+T|t+1}, \hat{d}_{t-L+1:t+T|t+1}^*, y_{t-L+1:t+1}). \tag{21}$$

On the other hand, from Assumption 3 at time $t+T$, with $d = \hat{d}_{t+T|t+1}^*$ and

$$x = \hat{x}_{t+T|t+1}^* \coloneqq \varphi(t+T; t-L+1, \hat{x}_{t-L+1|t+1}^*, \hat{u}_{t+1:t+T|t+1}^*, \hat{d}_{t-L+1:t+T|t+1}^*),$$

we conclude that there exists a control $\tilde{u}_{t+T} \in \mathcal{U}$ such that

$$\begin{aligned} q_{t+T+1}\big(f_{t+T}(\hat{x}_{t+T|t+1}^*, \tilde{u}_{t+T}, \hat{d}_{t+T|t+1}^*)\big) - q_{t+T}(\hat{x}_{t+T|t+1}^*) \\ + c_{t+T}(\hat{x}_{t+T|t+1}^*, \tilde{u}_{t+T}, \hat{d}_{t+T|t+1}^*) - \rho_{t+T}(\hat{d}_{t+T|t+1}^*) \leqslant 0. \end{aligned} \tag{22}$$

Moreover, we conclude from Assumption 2, that there exist vectors $\tilde{x}_{t-L}$, $\tilde{d}_{t-L} \in \mathcal{D}$, $\tilde{n}_{t-L} \in \mathcal{N}$ such that

$$\begin{aligned} \hat{x}_{t-L+1|t+1}^* &= f_{t-L}(\tilde{x}_{t-L}, u_{t-L}, \tilde{d}_{t-L}), \\ y_{t-L} &= g_{t-L}(\tilde{x}_{t-L}) + \tilde{n}_{t-L}, \end{aligned} \tag{23}$$

Using now (11a) in Assumption 1 at time t, we conclude that there also exists a finite scalar $J_t^* \in \mathbb{R}$ and a sequence $\hat{u}_{t:t+T-1|t}^* \in \mathcal{U}$ such that

$$J_t^* = \max_{\substack{\hat{x}_{t-L|t} \in \mathcal{X}, \\ \hat{d}_{t-L:t+T-1|t} \in \mathcal{D}}} J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}^*, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}). \tag{24}$$

Going back to (21), we then conclude that

$$J_{t+1}^* \leqslant J_{t+1}(\hat{x}_{t-L+1|t+1}^*, u_{t-L+1:t}, \hat{u}_{t+1:t+T-1|t}^*, \tilde{u}_{t+T}, \hat{d}_{t-L+1:t+T|t+1}^*, y_{t-L+1:t+1}) \tag{25}$$

because the minimization in (21) with respect to $\hat{u}_{t+1:t+T|t+1} \in \mathcal{U}$ must lead to a value no larger than what would be obtained by setting $\hat{u}_{t+1:t+T-1|t+1} = \hat{u}_{t+1:t+T-1|t}^*$ and $\hat{u}_{t+T|t+1} = \tilde{u}_{t+T}$.
Similarly, we can conclude from (24) that

$$\begin{aligned} J_t^* &\geqslant J_t(\tilde{x}_{t-L}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}^*, \tilde{d}_{t-L}, \hat{d}_{t-L+1:t+T-1|t+1}^*, y_{t-L:t}) \\ &= J_t(\tilde{x}_{t-L}, u_{t-L:t}, \hat{u}_{t+1:t+T-1|t}^*, \tilde{d}_{t-L}, \hat{d}_{t-L+1:t+T-1|t+1}^*, y_{t-L:t}), \end{aligned} \tag{26}$$

because the maximization in (24) with respect to $\hat{x}_{t-L|t}$ and $\hat{d}_{t-L:t+T-1|t}$ must lead to a value no smaller than what would be obtained by setting $\hat{x}_{t-L|t} = \tilde{x}_{t-L}$, $\hat{d}_{t-L|t} = \tilde{d}_{t-L}$ and $\hat{d}_{t-L+1:t+T-1|t} = \hat{d}^*_{t-L+1:t+T-1|t+1}$. The last equality in (26) is obtained by applying the control law (7).
Combining (25), (26), and (23) leads to

$$J^*_{t+1} - J^*_t \leqslant J_{t+1}\big(f_{t-L}(\tilde{x}_{t-L}, u_{t-L}, \tilde{d}_{t-L}), u_{t-L+1:t}, \hat{u}^*_{t+1:t+T-1|t}, \tilde{u}_{t+T}, \hat{d}^*_{t-L+1:t+T|t+1}, y_{t-L+1:t+1}\big) - J_t(\tilde{x}_{t-L}, u_{t-L:t}, \hat{u}^*_{t+1:t+T-1|t}, \tilde{d}_{t-L}, \hat{d}^*_{t-L+1:t+T-1|t+1}, y_{t-L:t}). \quad (27)$$

A crucial observation behind this inequality is that *both* terms $J_{t+1}(\cdot)$ and $J_t(\cdot)$ in the right-hand side of (27) are computed along a trajectory initialized at time $t-L$ with the *same initial state* $\tilde{x}_{t-L}$ and share the *same control input* sequence $u_{t-L:t}, \hat{u}^*_{t+1:t+T-1|t}$ and the *same disturbance input* sequence $\tilde{d}_{t-L}, \hat{d}^*_{t-L+1:t+T-1|t+1}$. We shall denote this common state trajectory by $\tilde{x}_s$, $s \in \{t-L, \dots, t+T\}$, and the shared control and disturbance sequences by

$$\tilde{d}_s \coloneqq \hat{d}^*_{s|t+1}, \quad \forall s \in \{t-L+1, \dots, t+T-1\},$$
$$\tilde{u}_s \coloneqq \begin{cases} u_s & s \in \{t-L, \dots, t\} \\ \hat{u}^*_{s|t} & s \in \{t+1, \dots, t+T-1\}. \end{cases}$$

The vectors $\tilde{u}_{t+T}$ and $\tilde{d}_{t-L}$ have been previously defined, but we now also define $\tilde{d}_{t+T} \coloneqq \hat{d}^*_{t+T|t+1}$, $\tilde{x}_{t+T+1} \coloneqq f_{t+T}(\tilde{x}_{t+T}, \tilde{u}_{t+T}, \tilde{d}_{t+T})$, and $\tilde{n}_s \coloneqq y_s - g_s(\tilde{x}_s)$, $s \in \{t-L, \dots, t\}$. All of these definitions enable us to express both terms $J_{t+1}(\cdot)$ and $J_t(\cdot)$ in the right-hand side of (27) as follows:

$$\begin{aligned} J^*_{t+1} - J^*_t &\leqslant \sum_{s=t+1}^{t+T} c_s(\tilde{x}_s, \tilde{u}_s, \tilde{d}_s) + q_{t+T+1}(\tilde{x}_{t+T+1}) - \sum_{s=t-L+1}^{t+1} \eta_s(\tilde{n}_s) - \sum_{s=t-L+1}^{t+T} \rho_s(\tilde{d}_s) \\ &\quad - \sum_{s=t}^{t+T-1} c_s(\tilde{x}_s, \tilde{u}_s, \tilde{d}_s) - q_{t+T}(\tilde{x}_{t+T}) + \sum_{s=t-L}^{t} \eta_s(\tilde{n}_s) + \sum_{s=t-L}^{t+T-1} \rho_s(\tilde{d}_s) \\ &= c_{t+T}(\tilde{x}_{t+T}, \tilde{u}_{t+T}, \tilde{d}_{t+T}) + q_{t+T+1}(\tilde{x}_{t+T+1}) - q_{t+T}(\tilde{x}_{t+T}) - \rho_{t+T}(\tilde{d}_{t+T}) + \eta_{t-L}(\tilde{n}_{t-L}) \\ &\quad + \rho_{t-L}(\tilde{d}_{t-L}) - c_t(\tilde{x}_t, \tilde{u}_t, \tilde{d}_t) - \eta_{t+1}(\tilde{n}_{t+1}). \end{aligned}$$

Equation (20) follows from this, (22), and the fact that $c_t(\cdot)$ and $\eta_{t+1}(\cdot)$ are both non-negative. ■

With most of the hard work done, we are now ready to prove the main result of this section.

Proof of Theorem 2. Using (11a) in Assumption 1, we conclude that

$$\begin{aligned} J^*_t &= \max_{\substack{\hat{x}_{t-L|t} \in \mathcal{X}, \\ \hat{d}_{t-L:t+T-1|t} \in \mathcal{D}}} J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}) \\ &\geqslant J_t(x_{t-L}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, d_{t-L:t}, 0_{t+1:t+T-1}, y_{t-L:t}) \\ &= J_t(x_{t-L}, u_{t-L:t}, \hat{u}^*_{t+1:t+T-1|t}, d_{t-L:t}, 0_{t+1:t+T-1}, y_{t-L:t}). \end{aligned}$$

The first inequality is a consequence of the fact that the maximum must lead to a value no smaller than what would have been obtained by setting $\hat{x}_{t-L|t}$ equal to the true state x_{t-L}, setting $\hat{d}_{t-L:t}$ equal to the

true (past) disturbances $d_{t-L:t}$ and setting $\hat{d}_{t+1:t+T-1}$ equal to zero. The final equality is obtained simply from the use of the control law (7).
To proceed, we replace $J_t(\cdot)$ by its definition in (10), while dropping all "future" positive terms in $c_s(\cdot)$, $s > t$ and $q_{t+T}(\cdot)$. This leads to

$$J_t^* \geqslant c_t(x_t, u_t, d_t) - \sum_{s=t-L}^{t} \eta_s(n_s) - \sum_{s=t-L}^{t} \rho_s(d_s). \tag{28}$$

Note that the future controls $\hat{u}^*_{t+1:t+T-1|t}$ disappeared because we dropped all the (positive) terms involving the value of the state past time t, and the summation over future disturbances also disappeared since we set all the future $\hat{d}_{t+1:t+T-1}$ to zero.
Adding both sides of (20) in Lemma 1 from time L to time $t-1$, leads to

$$J_t^* \leqslant J_L^* + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s) + \sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s), \quad \forall t \in \mathbb{Z}_{\geqslant L}. \tag{29}$$

The bound in (19) follows directly from (28) and (29). ■

4 Computation of Control by Solving a Pair of Coupled Optimizations

To implement the control law (7) we need to find the control sequence $\hat{u}^*_{t:\infty|t} \in \mathcal{U}$ that achieves the outer minimizations in (6) or (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:

$$-J_t^*(u_{t-L:t-1}, y_{t-L:t}) = \min_{\substack{\hat{x}_{t-L|t}\in\mathcal{X},\\ \hat{d}_{t-L:t+T-1|t}\in\mathcal{D}}} -J_t(\hat{x}_{t-L|t}, u_{t-L:t-1}, \hat{u}^*_{t:t+T-1|t}, \hat{d}_{t-L:t+T-1|t}, y_{t-L:t}) \tag{30a}$$

$$J_t^*(u_{t-L:t-1}, y_{t-L:t}) = \min_{\hat{u}_{t:t+T-1|t}\in\mathcal{U}} J_t(\hat{x}^*_{t-L|t}, u_{t-L:t-1}, \hat{u}_{t:t+T-1|t}, \hat{d}^*_{t-L:t+T-1|t}, y_{t-L:t}) \tag{30b}$$

(see, e.g., [14]). We introduced the "−" sign in (30a) 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 (30) can be re-written as

$$\begin{aligned} -J_t^*(u_{t-L:t-1}, y_{t-L:t}) = & \min_{(\hat{d}_{t-L:t+T-1|t}, \bar{x}_{t-L:t+T|t})\in\bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}]} \\ & - \sum_{s=t}^{t+T-1} c_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s) - q_{t+T}(\bar{x}_{t+T}) + \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\bar{x}_s)\big) + \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}_s) \end{aligned} \tag{31a}$$

$$\begin{aligned} J_t^*(u_{t-L:t-1}, y_{t-L:t}) = & \min_{(\hat{u}_{t:t+T-1|t}, \tilde{x}_{t-L+1:t+T|t})\in\bar{\mathcal{U}}[\hat{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}]} \\ & \sum_{s=t}^{t+T-1} c_s(\tilde{x}_s, \hat{u}_s, \hat{d}^*_s) + q_{t+T}(\tilde{x}_{t+T}) - \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\tilde{x}_s)\big) - \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}^*_s) \end{aligned} \tag{31b}$$

where

$$\bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}] \coloneqq \Big\{(\hat{d}_{t-L:t+T-1|t}, \bar{x}_{t-L:t+T|t}) : \hat{d}_{t-L:t+T-1|t} \in \mathcal{D},\ \bar{x}_{t-L:t+T|t} \in \mathcal{X},$$
$$\bar{x}_{s+1} = f_s(\bar{x}_s, u_s, \hat{d}_s), \forall s \in \{t-L, ..., t-1\},$$
$$\bar{x}_{s+1} = f_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s), \forall s \in \{t, ..., t+T-1\}\Big\}, \quad (32a)$$

$$\bar{\mathcal{U}}[\bar{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}] \coloneqq \Big\{(\hat{u}_{t:t+T-1|t}, \tilde{x}_{t-L+1:t+T|t}) : \hat{u}_{t:t+T-1|t} \in \mathcal{U},\ \tilde{x}_{t-L+1:t+T|t} \in \mathcal{X},$$
$$\tilde{x}_{t-L+1} = f_{t-L}(\bar{x}^*_{t-L}, u_{t-L}, \hat{d}^*_{t-L}),$$
$$\tilde{x}_{s+1} = f_s(\tilde{x}_s, u_s, \hat{d}^*_s), \forall s \in \{t-L+1, ..., t-1\},$$
$$\tilde{x}_{s+1} = f_s(\tilde{x}_s, \hat{u}_s, \hat{d}^*_s), \forall s \in \{t, ..., t+T-1\}\Big\}. \quad (32b)$$

Essentially, in each of the optimizations in (30) we introduced the values of the state from time 0 to time $t+T-1$ 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 $\varphi(t; 0, x_0, u_{0:t-1}, d_{0:t-1})$ that appears in the original optimizations (30) and that can be numerically poorly conditioned, e.g., for systems with unstable dynamics.

While the numerical method discussed in the next section can be used to solve either (30) or (31), all our numerical examples use the latter because it generally leads to simpler optimization problems and therefore we focus our discussion on that approach.

5 Interior-Point Method for Minimax Problems

The coupled optimizations in (30) or (31) can be viewed as a special case of the following more general problem: Find a pair $(u^*, d^*) \in \mathcal{U}[d^*] \times \mathcal{D}[u^*]$ that simultaneously solves the two coupled optimizations

$$f(u^*, d^*) = \min_{u \in \mathcal{U}[d^*]} f(u, d^*), \qquad g(u^*, d^*) = \min_{d \in \mathcal{D}[u^*]} g(u^*, d), \quad (33)$$

with

$$\mathcal{U}[d] \coloneqq \{u \in \mathbb{R}^{N_u} : F_u(u, d) \geqslant 0,\ G_u(u, d) = 0\}, \quad (34a)$$
$$\mathcal{D}[u] \coloneqq \{d \in \mathbb{R}^{N_d} : F_d(u, d) \geqslant 0,\ G_d(u, d) = 0\}, \quad (34b)$$

for given functions $f : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}$, $g : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}$, $F_u : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}^{M_u}$, $F_d : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}^{M_d}$, $G_u : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}^{K_u}$, $G_d : \mathbb{R}^{N_u} \times \mathbb{R}^{N_d} \to \mathbb{R}^{K_d}$. To map (31) to (33), one would associate the vectors $u \in \mathbb{R}^{N_u}$ and $d \in \mathbb{R}^{N_d}$ in (33) with the sequences in the sets $\bar{\mathcal{D}}[\cdot]$ and $\bar{\mathcal{U}}[\cdot]$ in (32). In this case, the equality constraints in (34) would typically correspond to the system dynamics and the inequality constraints in (34) would enforce that the state, control, and disturbance signals belong, respectively, to the sets $\mathcal{X}$, $\mathcal{U}$, and $\mathcal{D}$ introduced below (1).

Remark 6 (General optimization). The optimization in (33) is more general than the one in (31) in that the function being minimized in (31a) is the symmetric of the function being minimized in (31b), whereas in (33) f and g need not be the symmetric of each other. While this generalization does not appear to be particularly useful for our output-feedback model predictive control application, all the results that follow do apply to general functions f and g and can be useful for other applications. □

The following duality-like result provides the motivation for a primal-dual-like method to solve the coupled minimizations in (33). It provides a set of conditions, involving an unconstrained optimization, that provide an approximation to the solution of (33).

Lemma 2 (Approximate equilibrium)**.** *Suppose that we have found* primal variables $\hat{u} \in \mathbb{R}^{N_u}, \hat{d} \in \mathbb{R}^{N_d}$ *and* dual variables $\hat{\lambda}_{fu} \in \mathbb{R}^{M_u}, \hat{\lambda}_{gd} \in \mathbb{R}^{M_d}$, $\hat{\nu}_{fu} \in \mathbb{R}^{K_u}, \hat{\nu}_{gd} \in \mathbb{R}^{K_d}$ *that simultaneously satisfy all of the following conditions*[3]

$$G_u(\hat{u}, \hat{d}) = 0, \qquad G_d(\hat{u}, \hat{d}) = 0, \tag{35a}$$

$$\hat{\lambda}_{fu} \dot{\geqslant} 0, \qquad \hat{\lambda}_{gd} \dot{\geqslant} 0, \qquad F_u(\hat{u}, \hat{d}) \geqslant 0, \qquad F_d(\hat{u}, \hat{d}) \geqslant 0, \tag{35b}$$

$$L_f(\hat{u}, \hat{d}, \hat{\lambda}_{fu}, \hat{\nu}_{fu}) = \min_{u \in \mathbb{R}^{N_u}} L_f(u, \hat{d}, \hat{\lambda}_{fu}, \hat{\nu}_{fu}), \qquad L_g(\hat{u}, \hat{d}, \hat{\lambda}_{gd}, \hat{\nu}_{gd}) = \min_{d \in \mathbb{R}^{N_d}} L_g(\hat{u}, d, \hat{\lambda}_{gd}, \hat{\nu}_{gd}) \tag{35c}$$

where

$$\begin{aligned} L_f(u, d, \lambda_{fu}, \nu_{fu}) &\coloneqq f(u, d) - \lambda_{fu} F_u(u, d) + \nu_{fu} G_u(u, d), \\ L_g(u, d, \lambda_{gd}, \nu_{gd}) &\coloneqq g(u, d) - \lambda_{gd} F_d(u, d) + \nu_{gd} G_d(u, d), \qquad \forall u, d, \lambda, \nu. \end{aligned}$$

Then $(\hat{u}, \hat{d})$ *approximately satisfy* (33) *in the sense that*

$$f(\hat{u}, \hat{d}) \leqslant \epsilon_f + \min_{u \in \mathcal{U}[\hat{d}]} f(u, \hat{d}), \qquad g(\hat{u}, \hat{d}) \leqslant \epsilon_g + \min_{d \in \mathcal{D}[\hat{u}]} g(\hat{u}, d), \tag{36}$$

with

$$\epsilon_f \coloneqq \hat{\lambda}'_{fu} F_u(\hat{u}, \hat{d}), \qquad \epsilon_g \coloneqq \hat{\lambda}'_{gd} F_d(\hat{u}, \hat{d}). \quad \square \tag{37}$$

Proof of Lemma 2. The proof is a direct consequence of the following sequence of inequalities that start from the equalities in (35c) and use the conditions (35a)–(35b) and the definitions (37):

$$\begin{aligned} f(\hat{u}, \hat{d}) - \epsilon_f &= L_f(\hat{u}, \hat{d}, \hat{\lambda}_{fu}, \hat{\nu}_{fu}) - \hat{\nu}_{fu} G_u(\hat{u}, \hat{d}) = \min_{u \in \mathbb{R}^{N_u}} L_f(u, \hat{d}, \hat{\lambda}_{fu}, \hat{\nu}_{fu}) - 0 \\ &= \min_{u \in \mathbb{R}^{N_u}} f(u, \hat{d}) - \hat{\lambda}_{fu} F_u(u, \hat{d}) + \hat{\nu}_{fu} G_u(u, \hat{d}) \\ &\leqslant \max_{\lambda_{fu} \dot{\geqslant} 0, \nu_{fu}} \min_{u \in \mathbb{R}^{N_u}} f(u, \hat{d}) - \lambda_{fu} F_u(u, \hat{d}) + \nu_{fu} G_u(u, \hat{d}) \\ &\leqslant \max_{\lambda_{fu} \dot{\geqslant} 0, \nu_{fu}} \min_{u \in \mathcal{U}[\hat{d}]} f(u, \hat{d}) - \lambda_{fu} F_u(u, \hat{d}) + \nu_{fu} G_u(u, \hat{d}) \\ &= \min_{u \in \mathcal{U}[\hat{d}]} f(u, \hat{d}) \\ g(\hat{u}, \hat{d}) - \epsilon_g &= L_g(\hat{u}, \hat{d}, \hat{\lambda}_{gd}, \hat{\nu}_{gd}) - \hat{\nu}_{gd} G_d(\hat{u}, \hat{d}) = \min_{d \in \mathbb{R}^{N_d}} L_g(\hat{u}, d, \hat{\lambda}_{gd}, \hat{\nu}_{gd}) - 0 \\ &= \min_{d \in \mathbb{R}^{N_d}} g(\hat{u}, d) - \hat{\lambda}_{gd} F_d(\hat{u}, d) + \hat{\nu}_{gd} G_d(\hat{u}, d) \\ &\leqslant \max_{\lambda_{gd} \dot{\geqslant} 0, \nu_{gd}} \min_{d \in \mathbb{R}^{N_d}} g(\hat{u}, d) - \lambda_{gd} F_d(\hat{u}, d) + \nu_{gd} G_d(\hat{u}, d) \\ &\leqslant \max_{\lambda_{gd} \dot{\geqslant} 0, \nu_{gd}} \min_{d \in \mathcal{D}[\hat{u}]} g(\hat{u}, d) - \lambda_{gd} F_d(\hat{u}, d) + \nu_{gd} G_d(\hat{u}, d) \\ &= \min_{d \in \mathcal{D}[\hat{u}]} g(\hat{u}, d). \end{aligned}$$

□

[3] Given a vector $x \in \mathbb{R}^n$ and a scalar $a \in \mathbb{R}$, we denote by $x \dot{\geqslant} a$ the entry-wise comparison of x greater than or equal to a.

5.1 Interior-point primal-dual equilibria algorithm

The method proposed consists of using Newton iterations to solve a system of nonlinear equations on the primal variables $\hat{u} \in \mathbb{R}^{N_u}$, $\hat{d} \in \mathbb{R}^{N_d}$ and dual variables $\hat{\lambda}_{fu} \in \mathbb{R}^{M_u}$, $\hat{\lambda}_{gd} \in \mathbb{R}^{M_d}$, $\hat{\nu}_{fu} \in \mathbb{R}^{K_u}$, $\hat{\nu}_{gd} \in \mathbb{R}^{K_d}$ introduced in Lemma 2. The specific system of equations consists of:

1. the first-order optimality conditions for the unconstrained minimizations in (35c)[4]:

$$\nabla_u L_f(\hat{u}, \hat{d}, \hat{\lambda}_{fu}, \hat{\nu}_{fu}) = \mathbf{0}_{N_u}, \qquad \nabla_d L_g(\hat{u}, \hat{d}, \hat{\lambda}_{gd}, \hat{\nu}_{gd}) = \mathbf{0}_{N_d}; \tag{38}$$

2. the equality conditions (35a); and

3. the equations[5]

$$F_u(\hat{u}, \hat{d}) \odot \hat{\lambda}_{fu} = \mu \mathbf{1}_{M_u}, \qquad F_d(\hat{u}, \hat{d}) \odot \hat{\lambda}_{gd} = \mu \mathbf{1}_{M_d}, \tag{39}$$

for some $\mu > 0$, which leads to

$$\epsilon_f = M_u \mu, \qquad \epsilon_g = M_d \mu.$$

Since our goal is to find primal variables $\hat{u}, \hat{d}$ for which (36) holds with $\epsilon_f = \epsilon_g = 0$, we shall make the variable μ converge to zero as the Newton iterations progress. This is done in the context of an interior-point method, meaning that all variables will be initialized so that the inequality constraints (35b) hold and the progression along the Newton direction at each iteration will be selected so that these constraints are never violated.

The specific steps of the algorithm that follows are based on the primal-dual interior-point method for a single optimization, as described in [42]. To describe this algorithm, we define

$$z \coloneqq (\hat{u}, \hat{d}), \quad \lambda \coloneqq (\hat{\lambda}_{fu}, \hat{\lambda}_{gd}), \quad \nu \coloneqq (\hat{\nu}_{fu}, \hat{\nu}_{gd}), \quad G(z) \coloneqq \begin{bmatrix} G_u(\hat{u}, \hat{d}) \\ G_d(\hat{u}, \hat{d}) \end{bmatrix}, \quad F(z) \coloneqq \begin{bmatrix} F_u(\hat{u}, \hat{d}) \\ F_d(\hat{u}, \hat{d}) \end{bmatrix}, \tag{40}$$

which allow us to re-write (38), (35a), and (39) as

$$\nabla_u L_f(z, \lambda, \nu) = \mathbf{0}_{N_u}, \qquad \nabla_d L_g(z, \lambda, \nu) = \mathbf{0}_{N_d}, \qquad G(z) = \mathbf{0}_{K_u + K_d}, \qquad \lambda \odot F(z) = \mu \mathbf{1}_{M_u + M_d}, \tag{41a}$$

and (35b) as

$$\lambda \geqslant \mathbf{0}_{M_u + M_d}, \qquad F(z) \geqslant \mathbf{0}_{M_u + M_d}. \tag{41b}$$

Algorithm 1 (Primal-dual optimization)**.**

Step 1. *Start with estimates z_0, λ_0, ν_0 that satisfy the inequalities $\lambda_0 \geqslant 0$, $F(z_0) \geqslant 0$ in* (41b) *and set $\mu_0 = 1$ and $k = 0$. It is often a good idea to start with a value for z_0 that satisfies the equality constraint $G(z_0) = 0$, and $\lambda_0 = \mu_0 \mathbf{1}_{M_u + M_d} \oslash F(z_0)$, which guarantees that we initially have $\lambda_0 \odot F(z_0) = \mu_0 \mathbf{1}_{M_u + M_d}$.*

[4] Given an integer M, we denote by $\mathbf{0}_M$ and by $\mathbf{1}_M$ the M-vectors with all entries equal to 0 and 1, respectively.

[5] Given two vectors $x, y \in \mathbb{R}^n$ we denote by $x \odot y \in \mathbb{R}^n$ and by $x \oslash y \in \mathbb{R}^n$ the entry-wise product and division of the two vectors, respectively.

Step 2. *Linearize the equations in* (41a) *around a current estimate* z_k, λ_k, ν_k *and* μ_k*, leading to*

$$\begin{bmatrix} \nabla_{uz} L_f(z_k,\lambda_k,\nu_k) & \nabla_{u\nu} L_f(z_k) & \nabla_{u\lambda} L_f(z_k) \\ \nabla_{dz} L_g(z_k,\lambda_k,\nu_k) & \nabla_{d\nu} L_g(z_k) & \nabla_{d\lambda} L_g(z_k) \\ \nabla_z G(z_k) & 0 & 0 \\ \operatorname{diag}(\lambda_k)\nabla_z F(z_k) & 0 & \operatorname{diag}[F(z_k)] \end{bmatrix} \begin{bmatrix} \Delta z \\ \Delta \nu \\ \Delta \lambda \end{bmatrix} = - \begin{bmatrix} \nabla_u L_f(z_k,\lambda_k,\nu_k) \\ \nabla_d L_g(z_k,\lambda_k,\nu_k) \\ G(z_k) \\ F(z_k) \odot \lambda_k - \mu_k \mathbf{1} \end{bmatrix}. \tag{42}$$

Step 3. *Find the* affine scaling direction $[\,\Delta z_a' \;\Delta \nu_a' \;\Delta \lambda_a'\,]'$ *by solving* (42) *for* $\mu_k = 0$*:*

$$\begin{bmatrix} \nabla_{uz} L_f(z_k,\lambda_k,\nu_k) & \nabla_{u\nu} L_f(z_k) & \nabla_{u\lambda} L_f(z_k) \\ \nabla_{dz} L_g(z_k,\lambda_k,\nu_k) & \nabla_{d\nu} L_g(z_k) & \nabla_{d\lambda} L_g(z_k) \\ \nabla_z G(z_k) & 0 & 0 \\ \operatorname{diag}(\lambda_k)\nabla_z F(z_k) & 0 & \operatorname{diag}[F(z_k)] \end{bmatrix} \begin{bmatrix} \Delta z_a \\ \Delta \nu_a \\ \Delta \lambda_a \end{bmatrix} = - \begin{bmatrix} \nabla_u L_f(z_k,\lambda_k,\nu_k) \\ \nabla_d L_g(z_k,\lambda_k,\nu_k) \\ G(z_k) \\ F(z_k) \odot \lambda_k \end{bmatrix}.$$

Step 4. *Select scalings so that the inequalities in* (41b) *would not be violated along the affine scaling directions:*

$$\alpha_{\text{primal}} \coloneqq \max\big\{\alpha \in [0,1] : F(z_k + \alpha \Delta z_a) \geqslant 0\big\}, \qquad \alpha_{\text{dual}} \coloneqq \max\big\{\alpha \in [0,1] : \lambda_k + \alpha \Delta \lambda_a \geqslant 0\big\},$$

and define the following estimate for the "quality" of the affine scaling direction

$$\sigma \coloneqq \Big(\frac{F(z_k + \alpha_{\text{primal}} \Delta z_a) \odot (\lambda_k + \alpha_{\text{dual}} \Delta \lambda_a)}{F(z_k) \odot \lambda_k}\Big)^{\delta},$$

where δ *is a parameter typically selected equal to 3. The numerator* $F(z_k + \alpha_{\text{primal}} \Delta z_a) \odot (\lambda_k + \alpha_{\text{dual}} \Delta \lambda_a)$ *is the value one would obtain for* $\lambda \odot F(z)$ *by moving purely along the affine scaling directions. A small value for* σ *thus indicates that a significant reduction in* μ_k *is possible.*

Step 5. *Find the* search direction $[\,\Delta z_s' \;\Delta \nu_s' \;\Delta \lambda_s'\,]'$ *by solving* (42) *for* $\mu_k = \sigma \frac{F(z_k) \odot \lambda_k}{M_u + M_d}$*:*

$$\begin{bmatrix} \nabla_{uz} L_f(z_k,\lambda_k,\nu_k) & \nabla_{u\nu} L_f(z_k) & \nabla_{u\lambda} L_f(z_k) \\ \nabla_{dz} L_g(z_k,\lambda_k,\nu_k) & \nabla_{d\nu} L_g(z_k) & \nabla_{d\lambda} L_g(z_k) \\ \nabla_z G(z_k) & 0 & 0 \\ \operatorname{diag}(\lambda_k)\nabla_z F(z_k) & 0 & \operatorname{diag}[F(z_k)] \end{bmatrix} \begin{bmatrix} \Delta z_s \\ \Delta \nu_s \\ \Delta \lambda_s \end{bmatrix} = - \begin{bmatrix} \nabla_u L_f(z_k,\lambda_k,\nu_k) \\ \nabla_d L_g(z_k,\lambda_k,\nu_k) \\ G(z_k) \\ F(z_k) \odot \lambda_k - \sigma \frac{F(z_k) \odot \lambda_k}{M_u + M_d} \mathbf{1} \end{bmatrix}.$$

Step 6. *Update the estimates along the search direction so that the inequalities in* (41b) *hold strictly:*

$$z_{k+1} = z_k + \alpha_s \Delta z_s, \qquad \nu_{k+1} = \nu_k + \alpha_s \Delta \nu_s, \qquad \lambda_{k+1} = \lambda_k + \alpha_s \Delta \lambda_s,$$

where

$$\alpha_s \coloneqq \min\{\alpha_{\text{primal}}, \alpha_{\text{dual}}\},$$

and

$$\alpha_{\text{primal}} \coloneqq \max\big\{\alpha \in [0,1] : F(z_k + \frac{\alpha}{.99} \Delta z_s) \geqslant 0\big\}, \quad \alpha_{\text{dual}} \coloneqq \max\big\{\alpha \in [0,1] : \lambda_k + \frac{\alpha}{.99} \Delta \lambda_s \geqslant 0\big\}. \tag{43}$$

Also update μ_k *according to*

$$\mu_{k+1} = \xi \mu_k,$$

where the positive scalar ξ *is chosen such that* $\xi < 1$*.*

Step 7. *Repeat from Step 2 with an incremented value for k until*

$$\|\nabla_u L_f(z_k, \lambda_k, \nu_k)\| \leqslant \epsilon_u, \qquad \|\nabla_d L_g(z_k, \lambda_k, \nu_k)\| \leqslant \epsilon_d, \qquad \|G(z_k)\| \leqslant \epsilon_G, \qquad \lambda_k' F(z_k) \leqslant \epsilon_{\text{gap}}. \tag{44}$$

for sufficiently small tolerances ϵ_u, ϵ_d, ϵ_G, ϵ_{gap}. □

Remark 7 (Tuning parameters). The exponent δ that appears in Step 4 is typically selected equal to 3 but may be chosen differently depending on the nonconvexity and nonlinearity of the optimization problem being solved. Some tuning may be necessary to find a desirable value of δ. Similarly, the value of the parameter μ_k used in Step 5 also may need to be tuned for the particular optimization problem being solved especially in cases of nonconvexity and nonlinearity. For these cases, one can include the scalar ξ which multiplies μ_k to allow for tuning in different scenarios. □

When the functions L_f and L_g that appear in the unconstrained minimizations in (35c) have a single stationary point that corresponds to their global minimum, termination of the Algorithm 1 guarantees that the assumptions of Lemma 2 hold [up to the tolerances in (44)] and we obtain the desired solution to (33).

The desired uniqueness of the stationary point holds, e.g., when the function $f(u, d)$ is convex in u, $g(u, d)$ is convex in d, $F_u(u, d)$ is concave in u, $F_d(u, d)$ is concave in d, and $G_u(u, d)$ is linear in u, and $G_d(u, d)$ is linear in d. However, in practice the Algorithm 1 can find solutions to (33) even when these convexity assumptions do not hold. For problems for which one cannot be sure whether the Algorithm 1 terminated at a global minimum of the unconstrained problem, one may run several instances of the algorithm with random initial conditions. Consistent results for the optimizations across multiple initializations will provide an indication that a global minimum has been found.

Remark 8 (Smoothness). Algorithm 1 requires all the functions f, g, F_u, F_d, G_u, G_d to be twice differentiable for the computation of the matrices that appear in (42). However, this does not preclude the use of this algorithm in many problems where these functions are not differentiable because it is often possible to re-formulate non-smooth optimizations into smooth ones by appropriate transformations that often introduce additional optimization variables. Common examples of these transformations include the minimization of criteria involving ℓ_p norms, such as the "non-differentiable ℓ_1 optimization"

$$\min \left\{ \|A_{m\times n} x - b\|_{\ell_1} + \cdots : x \in \mathbb{R}^n, \dots \right\}$$

which is equivalent to the following "smooth optimization"

$$\min \left\{ v' \mathbf{1}_m + \cdots : x \in \mathbb{R}^n, v \in \mathbb{R}^m, -v \leqslant Ax - b \leqslant v, \dots \right\};$$

or the "non-differentiable ℓ_2 optimization"

$$\min \left\{ \|A_{m\times n} x - b\|_{\ell_2} + \cdots : x \in \mathbb{R}^n, \dots \right\}$$

$$\min \left\{ v + \cdots : x \in \mathbb{R}^n, v \geqslant 0, v^2 \geqslant (Ax - b)'(Ax - b), \dots \right\}.$$

More examples of such transformations can be found, e.g., in [49–52]. □

6 Validation through Simulation

In this Section we discuss several examples using the problem framework introduced in Section 2 and find solutions via simulation using the interior-point method described in Section 5.

For all of the following examples, we use a cost function of the form

$$J_t(x_{t-L}, u_{t-L:t+T-1}, d_{t-L:t+T-1}, y_{t-L:t}) = \sum_{s=t}^{t+T-1} \|x_s - r_s\|_2^2 + \lambda_u \sum_{s=t}^{t+T-1} \|u_s\|_2^2 - \lambda_n \sum_{s=t-L}^{t} \|n_s\|_2^2 - \lambda_d \sum_{s=t-L}^{t+T-1} \|d_s\|_2^2. \quad (45)$$

where r_s is a desired reference, and λ_u, λ_n, and λ_d are positive weighting constants.

Example 1 (Flexible beam). Consider a single-link flexible beam like the one described in [53], 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. Figure 3 shows an illustration of this example.

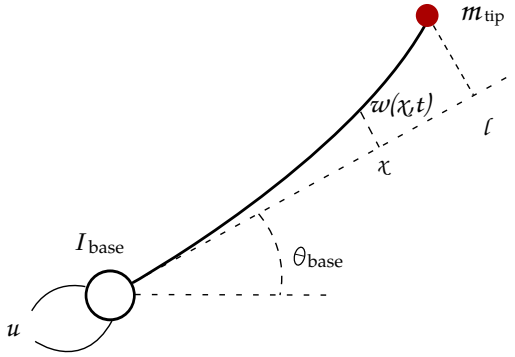

Figure 3: Illustration of the flexible beam.

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

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

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

$$C = \begin{bmatrix} 1.13 & 0.7225 & -0.2028 & 0.1220 & 0 & 0 & 0 & 0 \\ 1.0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1.0 & 0 & 0 & 0 \\ 0 & 0.9282 & -12.001 & -35.294 & 0 & 0 & 0 & 0 \end{bmatrix}. \quad (46)$$

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 optimization (9) with the cost function given in (45), with $\mathcal{U} := \{u_t \in \mathbb{R}^{n_u} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := \mathbb{R}^8$, and $\mathcal{D} := \{d_t \in \mathbb{R}^{n_d} : -d_{max} \leqslant d_t \leqslant d_{max}\}$. The numerical computation of solutions to the min-max optimization was performed using the primal-dual-like interior-point method described above.

The results depicted in Figure 4 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 $r_t := \alpha \operatorname{sgn}(\sin(\omega t))$ with $\alpha = 0.5$ and $\omega = 0.1$. The other parameters in the optimization have values $\lambda_u = 1$, $\lambda_d = 2$, $\lambda_n = 100$, $L = 5$, $T = 5$, $u_{max} = 0.8$, $d_{max} = 0.8$. 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. □

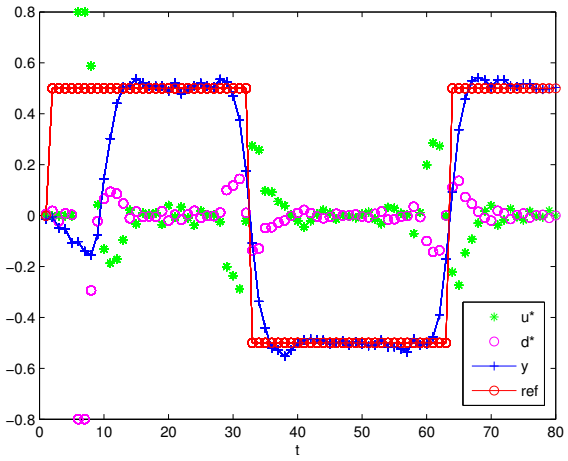

Figure 4: Simulation results of Flexible Beam example. The reference is in red, the measured output in blue, the control sequence in green, and the disturbance sequence in magenta.

Example 2 (Nonlinear Pursuit-Evasion). In this example we investigate a two-player pursuit-evasion game where the pursuer is modeled as a nonholonomic unicycle-type vehicle, and the evader is modeled as a double-integrator. The following equations are used to model this example

$$\begin{aligned} x^1_{t+1} &= x^1_t + v\cos(\theta_t), \\ x^2_{t+1} &= x^2_t + v\sin(\theta_t), \\ \theta_{t+1} &= \theta_t + u_t, \end{aligned} \qquad \begin{aligned} z^1_{t+1} &= z^1_t + d^1_t, \\ z^2_{t+1} &= z^2_t + d^2_t, \end{aligned}$$

where v is a constant scalar corresponding to the velocity of the pursuer, $x_t = \begin{bmatrix} x^1_t & x^2_t \end{bmatrix}' \in \mathbb{R}^2$ is the position of the pursuer at time t, $\theta_t \in [0, 2\pi]$ is the orientation of the pursuer at time t, $u_t \in \mathbb{R}$ is the control input at time t, $z_t = \begin{bmatrix} z^1_t & z^2_t \end{bmatrix}' \in \mathbb{R}^2$ is the position of the evader at time t, and $d_t = \begin{bmatrix} d^1_t & d^2_t \end{bmatrix}' \in \mathbb{R}^2$ is the evader's "speed" at time t. We assume that the control input u_t is bounded by the positive constant u_{max}, and "speed" of the evader d_t is bounded by the positive constant d_{max}. The output of the system is given by $y_t = \begin{bmatrix} x_t & z_t \end{bmatrix}' + \begin{bmatrix} n^1_t & n^2_t \end{bmatrix}'$, where $n_t = \begin{bmatrix} n^1_t & n^2_t \end{bmatrix}' \in \mathbb{R}^2$ is measurement noise.

The pursuer's goal is to make the distance between its position x_t and the position of the evader z_t as small as possible, so the pursuer wants to minimize the value of $\|z_t - x_t\|$. The evader's goal is to do the opposite, namely, maximize the value of $\|z_t - x_t\|$. The pursuer and evader try to achieve these goals by choosing appropriate values for u_t and d_t, respectively.

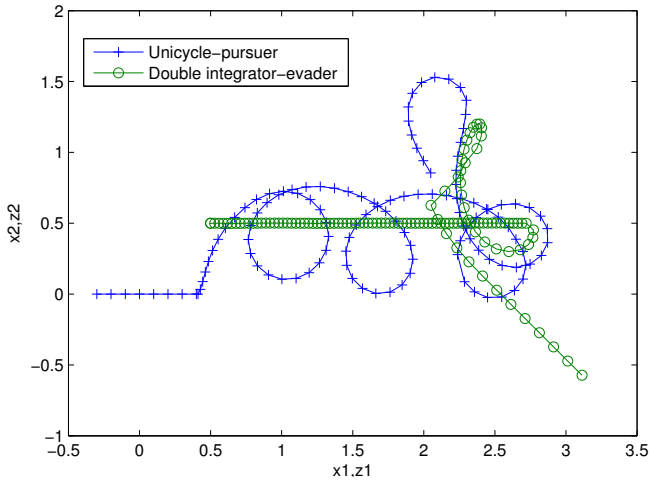

Figure 5: Results of nonlinear pursuit-evasion example. The path of the pursuer is shown with blue +'s, and the path of the evader is shown with green o's.

Figure 5 shows the results for this problem solving the optimization (9) with the cost function given in (45) with $\mathcal{U} := \{u_t \in \mathbb{R} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := \mathbb{R}^2 \times [0, 2\pi]$, and $\mathcal{D} := \{d_t \in \mathbb{R}^2 : -d_{max} \dot{\leqslant} d_t \dot{\leqslant} d_{max}\}$ and parameter values $L = 8$, $T = 12$, $v = .1$, $u_{max} = 0.5$, $d_{max} = 0.1$, $\lambda_u = 10$, $\lambda_d = 100$, and $\lambda_n = 10000$. The system evolved with zero control input until time $t = 8$, at which time the optimal control input (7) started to be applied. The value of d_t was selected to be constant $d_t = \begin{bmatrix} 0.03 & 0 \end{bmatrix}'$ until time $t = 75$ at which time the optimal $d^*_{t|t}$ obtained from the min-max optimization was applied. In this way, the evader moved at a constant fixed speed until time $t = 75$. After that time, the evader was "optimally" avoiding the pursuer by applying $d^*_{t|t}$. In the simulation shown in Figure 5, the noise process n_t was selected to be a zero-mean Gaussian independent and identically distributed random process. We see that, in this case, the evader is able to get away from the pursuer once it is playing "optimally". □

Example 3 (Model uncertainty). Consider a linear discrete time-invariant system given by

$$\begin{aligned} x_{t+1} &= \begin{bmatrix} 2 & -1 \\ 1 & a \end{bmatrix} x_t + \begin{bmatrix} 0.5 \\ 0 \end{bmatrix} u_t \\ y_t &= \begin{bmatrix} c & c \end{bmatrix} x_t \end{aligned}$$

where a and c are uncertain parameters known to belong in the intervals $a \in [-1, 1]$ and $c \in [0.25, 0.75]$. For this example, we consider solving an optimization of the form

$$\min_{u_{t:t+T-1|t} \in \mathcal{U}} \quad \max_{\substack{x_{t-L|t} \in \mathcal{X}, \\ a \in \mathcal{A}, \\ c \in \mathcal{C}}} \sum_{s=t}^{t+T-1} \|x_s\|_2^2 + \lambda_u \sum_{s=t}^{t+T-1} \|u_s\|_2^2 - \lambda_n \sum_{s=t-L}^{t} \|n_s\|_2^2, \tag{47}$$

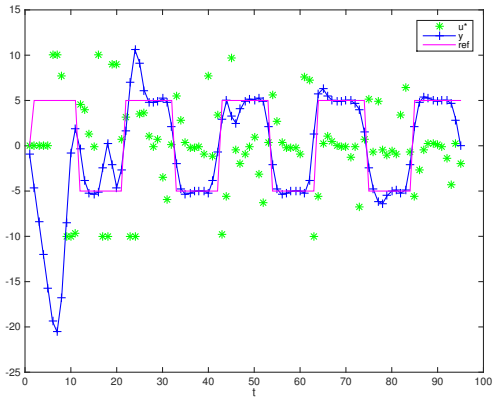

(a) Results of model uncertainty example. The measured output is in blue, the control sequence in green, and the reference signal in red.

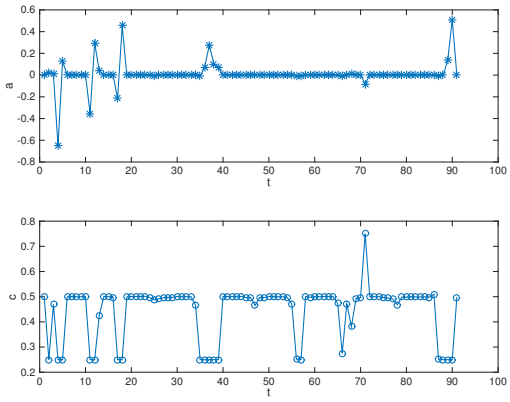

(b) Value of a and c from min-max optimization.

Figure 6: Results of model uncertainty example.

where $\mathcal{A} \coloneqq \{a \in \mathbb{R} : -1 \leqslant a \leqslant 1\}$ and $\mathcal{C} \coloneqq \{c \in \mathbb{R} : 0.25 \leqslant c \leqslant 0.75\}$, so we are optimizing over the worst-case values of the uncertain parameters a and c.

Figure 6 shows the results for this example solving the optimization (47) with $\mathcal{U} := \{u_t \in \mathbb{R} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := \mathbb{R}^2$, and parameter values $L = 5$, $T = 5$, $u_{max} = 10$, $\lambda_u = 0.1$, and $\lambda_n = 1000$. The reference signal is given by $r_t \coloneqq \alpha \operatorname{sgn}(\sin(\omega t))$ with $\alpha = 5$ and $\omega = 0.3$, and the true system has parameter values of $a = 0$ and $c = 0.5$. The system evolved with zero control input until time $t = 5$, at which time the optimal control input (7) started to be applied. In the simulation shown in Figure 6, the noise process n_t was selected to be a zero-mean Gaussian independent and identically distributed random process. We see that the controller is able to successfully regulate the system to the reference signal despite uncertain knowledge of the system.

Several more examples of using this combined control and estimation approach for parameter estimation can be found in [54].

Example 4 (Riderless Bicycle)**.** Consider a riderless bicycle as described in [55], where the control objective is to stabilize the bicycle in the upright position (i.e. zero roll angle). An approximate linearized discrete-time state-space model of the dynamics with a sampling time $T_s \coloneqq 0.1$ second is given by $x_{t+1} = Ax_t + B(u_t + d_t)$, $y_t = Cx_t + n_t$.

The state x_t of the system is comprised of the roll angle, the steering angle, and their respective derivatives. The control input u_t is the torque applied to the handlebars, d_t is an exogenous steer-torque disturbance, and n_t is measurement noise. The system matrices are given by

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 13.67 & 0.225 - 1.319 * v^2 & -0.164 * v & -0.552 * v \\ 4.857 & 10.81 - 1.125 * v^2 & 3.621 * v & -2.388 * v \end{bmatrix}, \quad B = \begin{bmatrix} 0 \\ 0 \\ -0.339 \\ 7.457 \end{bmatrix}, \quad C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \tag{48}$$

where v is the speed of the bicycle. The bicycle model is stable for speeds between $v = 3.4m/s$ and $v = 4.1m/s$ because the real part of the eigenvalues of the matrix A are negative. However, for lower speeds

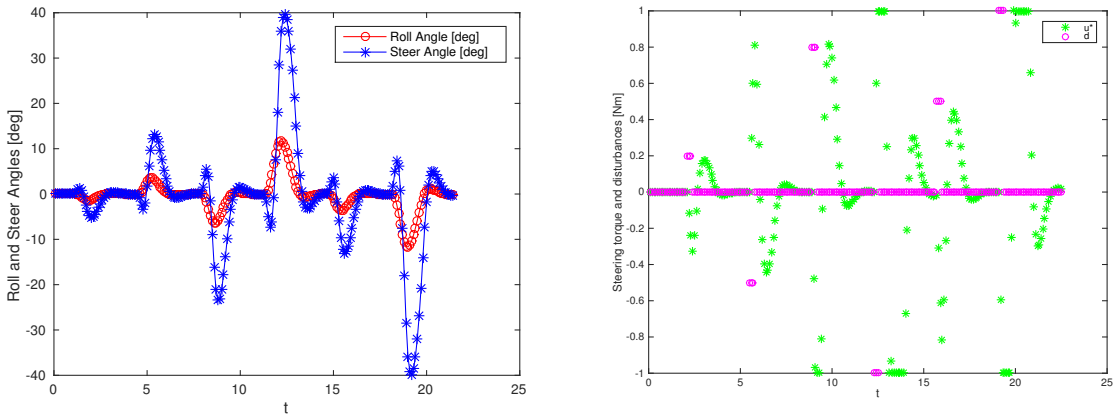

(a) Results of riderless bicycle example. The steer angle is in blue, and the roll angle is in red.

(b) Value of control inputs u^* (in blue) from min-max optimization and impulsive disturbances d (in red).

Figure 7: Results of riderless bicycle example.

or higher speeds up to about $10m/s$, the matrix A has at least one eigenvalue with negative real part, so the system is unstable. In this example, we assume that the bicycle is moving at a constant speed $v = 2m/s$.

Figure 7 shows the results for this example solving the optimization (9) with the cost function given in (45) with no zero reference signal, $\mathcal{U} := \{u_t \in \mathbb{R} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := [-\pi, \pi] \times [-\pi, \pi] \times \mathbb{R}^2$, and $\mathcal{D} := \{d_t \in \mathbb{R} : -d_{max} \leqslant d_t \leqslant d_{max}\}$ and parameter values $L = 10$, $T = 20$, $v = 2$, $u_{max} = 1$, $d_{max} = 1$, $\lambda_u = 0.01$, $\lambda_d = 100$, and $\lambda_n = 1000$. The system evolved with zero control input until time $t = 10$, at which time the optimal control input (7) started to be applied. The value of d_t was selected to be impulsive. In the simulation shown in Figure 7, the noise process n_t was selected to be a zero-mean Gaussian independent and identically distributed random process. We see that the controller is able to stabilize the riderless bicycle in the upright position (i.e. zero roll angle) in the presence of additive impulsive disturbances.

7 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 standard controllability and observability), 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 observability of the nonlinear dynamics and a terminal cost that is an ISS-control Lyapunov function with respect to the disturbance input.

Next we presented in Algorithm 1 a primal-dual-like interior-point method that can be used to solve the minimax optimization. We validated this algorithm by showing simulation results for both constrained linear and nonlinear examples. For convex problems, the global solution was found, and even for nonconvex problems, the algorithm converged to a solution.

Future work on these topics includes quantifying performance and proving convergence of Algorithm 1. The development of similar algorithms to solve these types of optimization problems and trade offs between methods could be investigated. For example, a barrier interior-point algorithm could be developed and compared to the primal-dual-like interior-point method given in Algorithm 1. We might expect the barrier

method to perform slower, but it could be more robust for non-convex poorly conditioned problems.

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] E. F. Camacho, C. Bordons, E. F. Camacho, and C. Bordons, *Model predictive control*, vol. 2. Springer London, 2004.

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

[4] L. Grüne and J. Pannek, *Nonlinear Model Predictive Control*. Springer, 2011.

[5] J. B. Rawlings, "Tutorial overview of model predictive control," *Control Systems, IEEE*, vol. 20, no. 3, pp. 38–52, 2000.

[6] S. J. Qin and T. A. Badgwell, "A survey of industrial model predictive control technology," *Control engineering practice*, vol. 11, no. 7, pp. 733–764, 2003.

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

[8] 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.

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

[10] 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.

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

[12] 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.

[13] 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.

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

[15] M. J. Tenny and J. B. Rawlings, "Efficient moving horizon estimation and nonlinear model predictive control," in *American Control Conference, 2002. Proceedings of the 2002*, vol. 6, pp. 4475–4480, IEEE, 2002.

[16] M. Diehl, H. J. Ferreau, and N. Haverbeke, "Efficient numerical methods for nonlinear MPC and moving horizon estimation," in *Nonlinear Model Predictive Control*, pp. 391–417, Springer, 2009.

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] C. V. Rao and J. B. Rawlings, "Nonlinear moving horizon state estimation," in *Nonlinear model predictive control*, pp. 45–69, Springer, 2000.

[25] C. V. Rao, J. B. Rawlings, and J. H. Lee, "Constrained linear state estimation-a moving horizon approach," *Automatica*, vol. 37, no. 10, pp. 1619–1628, 2001.

[26] 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.

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

[28] 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.

[29] 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.

[30] 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.

[31] 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.

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

[33] M. H. Wright, "The interior-point revolution in constrained optimization," in *High Performance Algorithms and Software in Nonlinear Optimization*, pp. 359–381, Springer, 1998.

[34] A. Forsgren, P. E. Gill, and M. H. Wright, "Interior methods for nonlinear optimization," *SIAM review*, vol. 44, no. 4, pp. 525–597, 2002.

[35] S. J. Wright, "Applying new optimization algorithms to model predictive control," in *AIChE Symposium Series*, vol. 93, pp. 147–155, Citeseer, 1997.

[36] C. V. Rao, S. J. Wright, and J. B. Rawlings, "Application of interior-point methods to model predictive control," *Journal of optimization theory and applications*, vol. 99, no. 3, pp. 723–757, 1998.

[37] L. Biegler and J. Rawlings, "Optimization approaches to nonlinear model predictive control," tech. rep., Argonne National Lab., IL (USA), 1991.

[38] L. T. Biegler, "Efficient solution of dynamic optimization and NMPC problems," in *Nonlinear model predictive control*, pp. 219–243, Springer, 2000.

[39] Y. Wang and S. Boyd, "Fast model predictive control using online optimization," *Control Systems Technology, IEEE Transactions on*, vol. 18, no. 2, pp. 267–278, 2010.

[40] D. M. de la Pena, T. Alamo, D. Ramirez, and E. Camacho, "Min-max model predictive control as a quadratic program," *Control Theory & Applications, IET*, vol. 1, no. 1, pp. 328–333, 2007.

[41] M. Diehl and J. Bjornberg, "Robust dynamic programming for min-max model predictive control of constrained uncertain systems," *Automatic Control, IEEE Transactions on*, vol. 49, no. 12, pp. 2253–2257, 2004.

[42] L. Vandenberghe, "The CVXOPT linear and quadratic cone program solvers," tech. rep., Univ. California, Los Angeles, 2010.

[43] 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.

[44] 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.

[45] D. A. Copp and J. P. Hespanha, "Conditions for saddle-point equilibria in output-feedback MPC with MHE," in *2016 American Control Conference (ACC)*, pp. 13–19, IEEE, July 2016.

[46] 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.

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

[48] 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.

[49] M. Grant and S. Boyd, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control* (V. Blondel, S. Boyd, and H. Kimura, eds.), vol. 371 of *Lecture Notes in Control and Information Sciences*, pp. 95–110, Springer Berlin / Heidelberg, 2008.

[50] F. Facchinei, H. Jiang, and L. Qi, "A smoothing method for mathematical programs with equilibrium constraints," *Mathematical programming*, vol. 85, no. 1, pp. 107–134, 1999.

[51] J. Kreimer and R. Y. Rubinstein, "Nondifferentiable optimization via smooth approximation: General analytical approach," *Annals of Operations Research*, vol. 39, no. 1, pp. 97–119, 1992.

[52] D. Y. Gao, "Canonical dual transformation method and generalized triality theory in nonsmooth global optimization*," *Journal of Global Optimization*, vol. 17, no. 1-4, pp. 127–160, 2000.

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

[54] D. A. Copp and J. P. Hespanha, "Addressing adaptation and learning in the context of model predictive control with moving horizon estimation," in *Control of Complex Systems: Theory and Applications* (K. G. Vamvoudakis and S. Jagannathan, eds.), ch. 6, Elsevier, July 2016.

[55] V. Cerone, D. Andreo, M. Larsson, and D. Regruto, "Stabilization of a riderless bicycle," *IEEE Control Systems Magazine*, vol. 30, pp. 23–32, October 2010.