Nonlinear Output-Feedback Model Predictive Control with Moving Horizon Estimation: Illustrative Examples

David A. Copp and João P. Hespanha*

October 28, 2015

Abstract

We review a recently introduced method to efficiently solve online optimization problems that appear in output-feedback model predictive control (MPC) and moving-horizon state estimation (MHE). The novel feature of this approach is that it solves both the MPC and MHE problems simultaneously as 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 give results showing that the state of the system remains bounded and establish bounds on the tracking error for trajectory tracking problems.

The min-max optimization that arises in our approach can be solved using a primal-dual-like interior-point method, developed especially for this purpose. 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. For all these examples, the interior-point method solver takes on average less than 6 ms to compute the control signal on a regular laptop computer.

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 useful 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

*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

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, 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].

Other important considerations when implementing an MPC controller include robustness to model parameter uncertainty, additive 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 explicitly handles constraints and computes the optimal current estimate of the state by solving an online optimization problem over a fixed number of past measurements. Therefore, the computational cost does not grow as more measurements become available. Nonlinear MPC and MHE are are both discussed in [12]. A useful overview of constrained nonlinear moving horizon state estimation is given in [13], and more recent results regarding stability of MHE can be found in [14].

Thus far, results on the stability of output-feedback control schemes based on MPC and MHE (especially for nonlinear systems) are limited. In this report, we consider the output-feedback of nonlinear systems with uncertainty and disturbances, and formulate the MPC problem as a min-max optimization as introduced in [15, 16]. 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 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 our 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 [17] and presumes appropriate forms of observability and controllability of the closed-loop system. We additionally require the minor assumptions that the dynamics are reversible and that there exists a terminal cost that is an ISS-control Lyapunov function with respect to a disturbance input.

Several algorithms are available to numerically solve the class of min-max optimization problems that we discuss here. A few methods are discussed in [18] and [19] and include sequential quadratic programming, interior-point methods, and others. We use a primal-dual-like interior-point algorithm as described in [16] to solve this min-max optimization. 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.

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 5, and we provide some conclusions and directions for future research in Section 6.

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 [20] and [21], and a game theoretic approach for robust constrained nonlinear MPC is proposed in [22]. Recent studies of input-to-state stability of min-max MPC can be found in [23–25], however these references do not investigate the use of output-feedback. Nominal or inherent robustness of MPC has also been studied in [3, 26].

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, 27, 28]. Recent results regarding stability of MHE can be found in [14]. Nonlinear MPC and MHE are both discussed in [12]. Some joint stability results for state estimation and control are given in [29], 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 [30]. Results on robust output-feedback MPC for constrained linear systems can be found in [31] using a state observer for estimation, and in [32] using MHE for estimation. Fewer results are available for nonlinear output-feedback MPC, although notable exceptions are [3, 33].

There is extensive literature involving the derivation of methods to numerically solve optimization problems, such as those that appear in MPC and MHE, efficiently and reliably. A good resource for studying convex optimization problems and methods to solve them (including interior-point and primal-dual interior-point methods) is the book [34]. Work regarding interior-point methods can be found in [35, 36] and primal-dual interior-point methods in particular in [37]. The application of interior-point algorithms as a method to solve MPC problems is discussed in [38]. Other early work on efficient numerical methods for solving MPC problems are given in [39], [40], and [37]. Advances in computational efficiency have allowed for the fast solution of MPC problems using online optimization such as in the more recent work [41]. The real-time solution of the MHE problem for small dimensional nonlinear models is given in [18]. Considering specifically numerical methods for min-max MPC optimization problems, the authors in [42] 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 [43], while sequential quadratic programming and interior-point methods for solving nonlinear MPC with MHE problems are discussed in [19]. The particular primal-dual-like interior-point method that we implement is inspired by the primal-dual interior-point method for a single optimization given in [44].

2 Problem Formulation

As in [15] and [16], we consider the control of a time-varying nonlinear discrete-time process of the form

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

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 $\mathcal{D} \subset \mathbb{R}^{n_d}$,

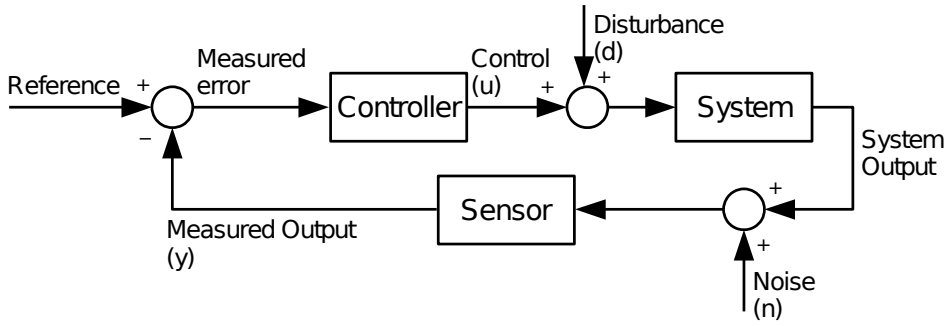

Figure 1: Block diagram depicting the problem formulation.

and the *measurement noise* $n_t \in \mathbb{R}^{n_n}$. 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 $\sum_{t=0}^{\infty} c_t(x_t, u_t, d_t) \leqslant \gamma + \sum_{t=0}^{\infty} \eta_t(n_t) + \sum_{t=0}^{\infty} \rho_t(d_t)$.

Remark 1 (Quadratic case). While the results presented here are very general in what regards the criterion (2), 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. However, the criterion (2) is compatible with much more sophisticated penalty functions that include, e.g., the penalty functions used in economic MPC in which the operating cost of the plant is used directly in the objective function [45]. □

2.1 Finite-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

$$\sum_{s=t}^{t+T-1} c_s(x_s, u_s, d_s) + q_{t+T}(x_{t+T}) - \sum_{s=t-L}^{t} \eta_s(n_s) - \sum_{s=t-L}^{t+T-1} \rho_s(d_s), \tag{3}$$

under worst-case assumptions on the *unknown* system's initial condition x_{t-L}, 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. This finite-horizon 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. This optimization criterion

also only contains $L+1 \in \mathbb{Z}_{\geqslant 1}$ terms of the measurement cost $\eta_s(n_s)$. The function $q_{t+T}(x_{t+T})$ acts as a terminal cost in order to penalize the "final" state at time $t+T$.

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 (3) 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 (3) 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.

At a given time $t \in \mathbb{Z}_{\geqslant 0}$, we do not know the value of the variables x_{t-L} and $d_{t-L:t+T-1}$ [1] on which the value of the criterion (3) depends, so we optimize for the future controls $u_{t:t+T-1}$ under worst-case assumptions on x_{t-L} and $d_{t-L:t+T-1}$, leading to the following min-max optimization

$$J_t^* = \min_{\hat{u}_{t:t+T-1} \in \mathcal{U}} \quad \max_{\hat{x}_{t-L} \in \mathcal{X}, \hat{d}_{t-L:t+T-1} \in \mathcal{D}} \sum_{s=t}^{t+T-1} c_s(\hat{x}_s, \hat{u}_s, \hat{d}_s) + q_{t+T}(\hat{x}_{t+T}) \\ - \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\hat{x}_s)\big) - \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}_s), \quad (4)$$

with the understanding that

$$\hat{x}_{s+1} = \begin{cases} f_s(\hat{x}_s, u_s, \hat{d}_s) & \text{for } t-L \leqslant s < t, \\ f_s(\hat{x}_s, \hat{u}_s, \hat{d}_s) & \text{for } t \leqslant s < t+T. \end{cases}$$

We can view the optimization variables $\hat{x}_{t-L}$ and $\hat{d}_{t-L:t+T-1}$ as (worst-case) estimates of the initial state and disturbances, respectively, based on the past inputs $u_{t-L:t-1}$ and outputs $y_{t-L:t}$ 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

$$\hat{u}_{t:t+T-1}^* = \{\hat{u}_t^*, \hat{u}_{t+1}^*, \hat{u}_{t+2}^*, \ldots, \hat{u}_{t+T-1}^*\} \in \mathcal{U}$$

that minimizes (4), leading to the following control law:

$$u_t = \hat{u}_t^*, \quad \forall t \geqslant 0. \quad (5)$$

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

2.2 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 (5) that minimizes (4) can also be determined by the optimization

$$\min_{\hat{u}_{t:t+T-1} \in \mathcal{U}} \quad \max_{\hat{d}_{t:t+T-1} \in \mathcal{D}} \sum_{s=t}^{t+T-1} c_s(\hat{x}_s, \hat{u}_s, \hat{d}_s) + q_{t+T}(\hat{x}_{t+T}) - \sum_{s=t}^{t+T-1} \rho_s(\hat{d}_s),$$

which is essentially the robust model predictive control with terminal cost considered in [22, 46].

[1] Here and below, we use the notation $x_{t_1:t_2}$ to denote the time series $x_{t_1}, x_{t_1+1}, \ldots, x_{t_2-1}, x_{t_2}$.

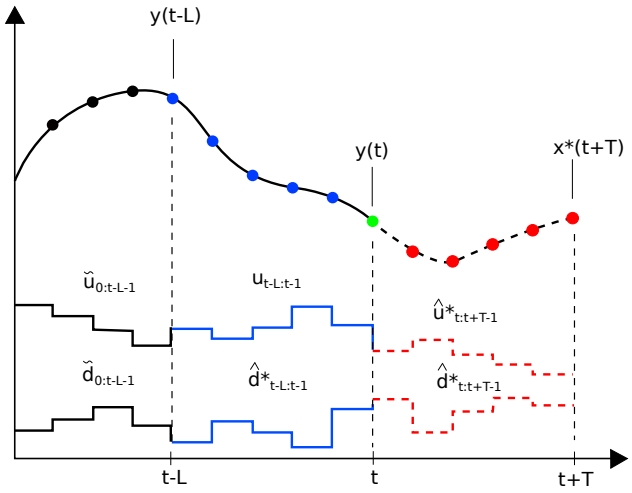

Figure 2: Finite-horizon control and estimation problems. The elements in blue correspond to the MHE problem, and the elements in red correspond to the MPC problem.

2.3 Relationship with Moving-Horizon Estimation

When setting both $c_s(\cdot)$ and $q_{t+T}(\cdot)$ equal to zero in the criterion (4), this optimization no longer depends on $u_{t:t+T-1}$ and $d_{t:t+T-1}$, so the optimization in (4) simply becomes

$$\max_{\hat{x}_{t-L}\in\mathcal{X},\hat{d}_{t-L:t-1}\in\mathcal{D}} -\sum_{s=t-L}^{t}\eta_s\big(y_s-g_s(\hat{x}_s)\big)-\sum_{s=t-L}^{t-1}\rho_s(\hat{d}_s),$$

where now the optimization criterion only contains a finite number of terms that recedes as the current time t advances, which is essentially the moving horizon estimation problem considered in [13, 14].

3 Main Results

We now show that for the finite-horizon optimization introduced in Section 2.1, the control law (5) 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 (5) is that the outer minimization in (4) leads to a finite value for the optimum. However, for the stability results in this section we actually ask for the existence of a saddle-point solution to the min-max optimization in (4), which is a common requirement in game theoretical approaches to control design [17]:

Assumption 1 (Saddle-point). The min-max optimization (4) always has a finite-valued saddle-point solution for which the min and max commute. □

Assumption 1 is described in more detail in [15] and [16] and presumes an appropriate form of *observability/detectability* adapted to the criterion $\sum_{s=t}^{t+T-1} c_s(x_s,u_s,d_s)$. In particular, it implies that we can bound the size of the *current* state using past outputs and past/future input disturbances.

To ensure controllability and to establish state boundedness under the control (5) defined by the *finite-horizon* optimization criterion (4), one needs additional assumptions regarding the dynamics and the terminal cost $q_t(\cdot)$.

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

$$x_{t+1} = f_t(\tilde{x}_t, u_t, \tilde{d}_t). \tag{6}$$

□

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

$$q_{t+1}\big(f_t(x, u, d)\big) - q_t(x) \leqslant -c_t(x, u, d) + \rho_t(d), \quad \forall d \in \mathcal{D}. \tag{7}$$

□

Assumption 2 is very mild and essentially means that the sets of disturbances $\mathcal{D}$ and past states $\mathcal{X}$ are sufficiently rich to allow for a jump to any future state in $\mathcal{X}$. Assumption 3 plays the role of a common assumption in model predictive control, namely that the terminal cost must be a control Lyapunov function for the closed-loop [47]. In the absence of the disturbance d_t, (7) 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 [48]. 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 [49].

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

3.1 Finite-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 *finite-horizon* criterion (3). This result first appeared in [15], and its proof can be found in [16].

Theorem 1 (Finite-horizon cost-to-go bound)**.** *Suppose that Assumptions 1, 2, and 3 hold. Then there exists a finite constant J_0^* and vectors $\tilde{d}_s \in \mathbb{R}^{n_d}$, $\tilde{n}_s \in \mathbb{R}^{n_n}$, $\forall s \in \{0, 1, \ldots, t-L-1\}$ for which*

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

and the trajectories of the process (1) *with control* (5) *defined by the finite-horizon optimization* (4) *satisfy*

$$c_t(x_t, u_t, d_t) \leqslant J_0^* + \sum_{s=0}^{t-L-1} \eta_s(\tilde{n}_s) + \sum_{s=0}^{t-L-1} \rho_s(\tilde{d}_s) + \sum_{s=t-L}^{t} \eta_s(n_s) + \sum_{s=t-L}^{t} \rho_s(d_s), \quad \forall t \in \mathbb{Z}_{\geqslant 0} \tag{8}$$

□

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 (8) 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]. In what follows, for simplicity of presentation we no longer distinguish between $\tilde{d}_s$ and d_s or $\tilde{n}_s$ and n_s, but instead simply write d_s and n_s.

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

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

we conclude from (8) that, along trajectories of the closed-loop system, we have

$$\alpha(\|x_t\|) \leqslant J_0^* + \sum_{s=0}^{t} \beta(\|n_s\|) + \sum_{s=0}^{t} \delta(\|d_s\|), \quad \forall t \in \mathbb{Z}_{\geqslant 0}. \tag{9}$$

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

$$c_t(x,u,d) \geqslant \lambda^{-t}\alpha(\|x\|), \quad \eta_t(n) \leqslant \lambda^{-t}\beta(\|n\|), \quad \rho_t(d) \leqslant \lambda^{-t}\delta(\|d\|), \quad \forall x \in \mathbb{R}^{n_x}, u \in \mathbb{R}^{n_u}, d \in \mathbb{R}^{n_d}, n \in \mathbb{R}^{n_n}, \tag{10}$$

for some $\lambda \in (0,1)$, in which case we conclude from (8) that

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

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* (10) *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 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* (5) *defined by the finite-horizon optimization* (4)*. 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 (10) typically arise from criterion of the form

$$\sum_{s=t}^{t+T-1} \lambda^{-s} c(x_s, u_s, d_s) - \sum_{s=t-L}^{t} \lambda^{-s}\eta(n_s) - \sum_{s=t-L}^{t+T-1} \lambda^{-s}\rho(d_s)$$

that weight the future more than the past. In this case, (10) holds for functions α, β, and δ such that $c(x,u,d) \geqslant \alpha(\|x\|)$, $\eta(n) \leqslant \beta(\|n\|)$, and $\rho(d) \leqslant \delta(\|d\|)$, $\forall x, u, d, n$. □

[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.

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

$$\sum_{s=t}^{t+T-1} \lambda^{-s} c(x_s - z_s, u_s, d_s) - \sum_{s=t-L}^{t} \lambda^{-s} \eta(n_s) - \sum_{s=t-L}^{t+T-1} \lambda^{-s} \rho(d_s).$$

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 (8) that

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

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.

4 Computation of Control by Solving a Pair of Coupled Optimizations

To implement the control law (5) we need to find the control sequence $\hat{u}^*_{t:t+T-1} \in \mathcal{U}$ that achieves the outer minimizations in (4). In view of Assumption 1, the desired control sequence must be part of a saddle-point. From the perspective of numerically computing this saddle point, it is convenient to use the following characterization of the saddle point:

$$-J_t^* = \min_{(\hat{d}_{t-L:t+T-1}, \bar{x}_{t-L:t+T}) \in \bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}]} - \sum_{s=t}^{t+T-1} c_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s) - q_{t+T}(\bar{x}_{t+T}) + \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\bar{x}_s)\big) + \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}_s) \tag{11a}$$

$$J_t^* = \min_{(\hat{u}_{t-L:t+T-1}, \tilde{x}_{t-L:t+T}) \in \bar{\mathcal{U}}[\bar{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}]} \sum_{s=t}^{t+T-1} c_s(\tilde{x}_s, \hat{u}_s, \hat{d}^*_s) + q_{t+T}(\tilde{x}_{t+T}) - \sum_{s=t-L}^{t} \eta_s\big(y_s - g_s(\tilde{x}_s)\big) - \sum_{s=t-L}^{t+T-1} \rho_s(\hat{d}^*_s) \tag{11b}$$

where

$$\bar{\mathcal{D}}[u_{t-L:t-1}, \hat{u}^*_{t:t+T-1}] \coloneqq \Big\{ (\hat{d}_{t-L:t+T-1|t}, \bar{x}_{t-L:t+T}) : \hat{d}_{t-L:t+T-1} \in \mathcal{D},\ \bar{x}_{t-L:t+T} \in \mathcal{X}, \\ \bar{x}_{s+1} = f_s(\bar{x}_s, u_s, \hat{d}_s), \forall s \in \{t-L, \dots, t-1\}, \\ \bar{x}_{s+1} = f_s(\bar{x}_s, \hat{u}^*_s, \hat{d}_s), \forall s \in \{t, \dots, t+T-1\} \Big\}, \tag{12a}$$

$$\bar{\mathcal{U}}[\bar{x}^*_{t-L}, \hat{d}^*_{t-L:t+T-1}] \coloneqq \Big\{ (\hat{u}_{t:t+T-1}, \tilde{x}_{t-L+1:t+T}) : \hat{u}_{t:t+T-1} \in \mathcal{U},\ \tilde{x}_{t-L+1:t+T} \in \mathcal{X}, \\ \tilde{x}_{t-L+1} = f_s(\bar{x}^*_{t-L}, u_{t-L}, \hat{d}^*_{t-L}),$$

$$\tilde{x}_{s+1} = f_s(\tilde{x}_s, u_s, \hat{d}_s^*), \forall s \in \{t-L+1, ..., t-1\},$$
$$\tilde{x}_{s+1} = f_s(\tilde{x}_s, \hat{u}_s, \hat{d}_s^*), \forall s \in \{t, ..., t+T-1\}\Big\}. \quad (12b)$$

Essentially, we introduced the values of the state from time t-L+1 to time $t+T$ as additional optimization variables that are constrained by the system dynamics.

Using this characterization of the saddle point, we have developed a primal-dual-like interior-point method to find it, inspired by the primal-dual interior-point method for a single optimization given in [44]. We do not discuss this method here, but details of the method can be found in [16]. 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.

5 Validation through Simulation

In this section we consider several examples and present closed-loop simulations using the control approach introduced in Section 2. For all of the following examples, we use a cost function of the form

$$\sum_{s=t}^{t+T-1} \|h_s(x_s)\|_2^2 + \lambda_u \sum_{s=t}^{t+T-1} \|u_s\|_2^2 - \lambda_n \sum_{s=t-L}^{t} \|n_s\|_2^2 - \lambda_d \sum_{s=t-L}^{t+T-1} \|d_s\|_2^2. \quad (13)$$

where $h_s(x_s)$ is a function of the state x_s that is especially relevant for the example under consideration, and λ_u, λ_n, and λ_d are positive weighting constants.

Example 1 (Flexible beam)**.** In this example we consider a single-link flexible beam like the one described in [50]. This system is 8-dimensional and exhibits oscillatory modes that are very lightly damped . An illustration of this system is given in Figure 3. 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.

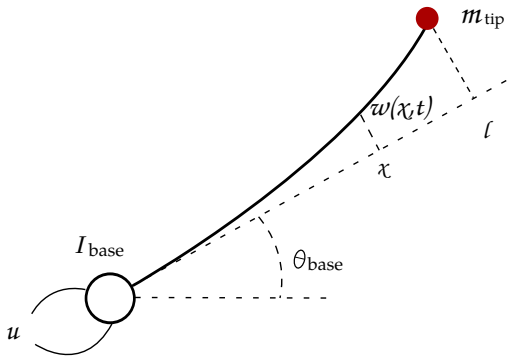

Figure 3: Illustration of the flexible beam.

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

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

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

$$C = \begin{bmatrix} 1.13 & 0.7225 & -0.2028 & 0.1220 & 0 & 0 & 0 & 0 \\ 1.0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1.0 & 0 & 0 & 0 \\ 0 & 0.9282 & -12.001 & -35.294 & 0 & 0 & 0 & 0 \end{bmatrix}. \tag{14}$$

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 (4) with the cost function given in (13), with $h_s(x_s) = p_s - r_s$, where p_s is the position of the mass at the tip of the beam and r_s is a desired reference trajectory, $\mathcal{U} \coloneqq \{u_t \in \mathbb{R}^{n_u} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} \coloneqq \mathbb{R}^8$, and $\mathcal{D} \coloneqq \{d_t \in \mathbb{R}^{n_d} : -d_{max} \leqslant d_t \leqslant d_{max}\}$.

The results depicted in Figure 4 show the response of the closed loop system under the control law (5) when our goal is to regulate the mass at the tip of the beam to a desired reference $r_t \coloneqq \alpha \operatorname{sgn}(\sin(\omega t))$ with $\alpha = 0.5$ and $\omega = 0.1$. The other parameters in the optimization have values $\lambda_u = 1$, $\lambda_d = 2$, $\lambda_n = 100$, $L = 5$, $T = 5$, $u_{max} = 0.8$, $d_{max} = 0.8$. 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 (5) started to be applied along with the optimal worst-case disturbance d_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 a standard deviation of 0.01. □

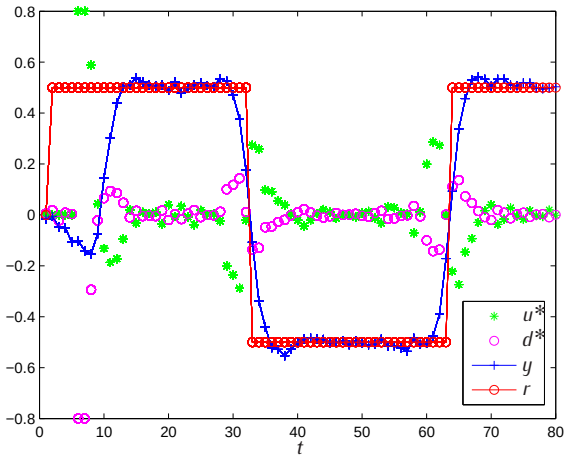

Figure 4: Simulation results of Flexible Beam example. The reference is in red, the measured output in blue, the control sequence in green, and the disturbance sequence in magenta.

Example 2 (Riderless bicycle)**.** In this example we consider an unstable linear system with impulsive additive input disturbances. It corresponds to the riderless bicycle described in [51], where the control objective is to

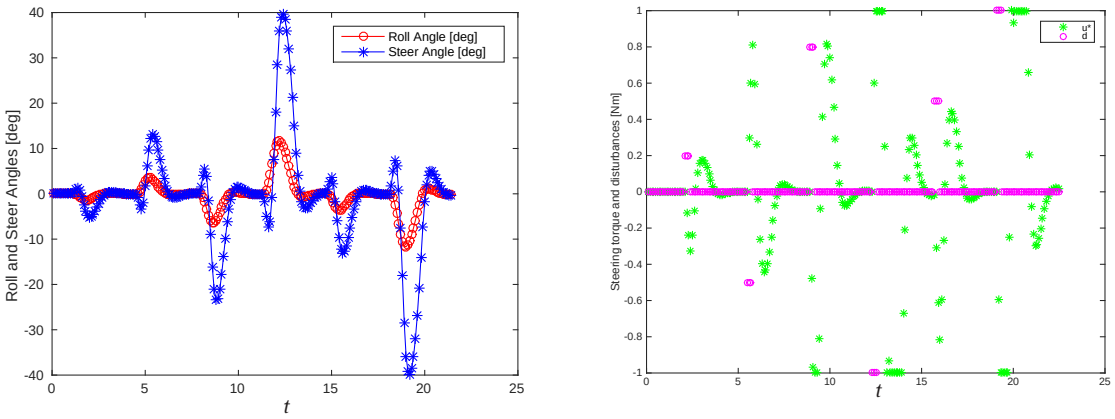

(a) Results of riderless bicycle example. The steer angle is in blue, and the roll angle is in red.

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

Figure 5: Results of riderless bicycle example.

stabilize the bicycle in the upright position (i.e. zero roll angle). An approximate linearized continuous-time state-space model is given by $\dot{x} = Ax + B(u + d)$, $y = Cx + n$.

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 steering-torque disturbance, and n_t is measurement noise. The system matrices are given by

$$
\begin{aligned}
A &= \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 13.67 & 0.225 - 1.319 * v^2 & -0.164 * v & -0.552 * v \\ 4.857 & 10.81 - 1.125 * v^2 & 3.621 * v & -2.388 * v \end{bmatrix}, \\
B &= \begin{bmatrix} 0 & 0 & -0.339 & 7.457 \end{bmatrix}', \\
C &= \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},
\end{aligned} \tag{15}
$$

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 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 $v = 2m/s$ and discretize the continuous-time model using a zero-order hold with a sampling time of $T_s := 0.1$ seconds.

Figure 5 shows the results for this example solving the optimization (4) with the cost function given in (13) with $h_s(x_s)$ equal to the roll angle of the bicycle, $\mathcal{U} := \{u_t \in \mathbb{R} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := [-\pi, \pi] \times [-\pi, \pi] \times \mathbb{R}^2$, 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 (5) started to be applied. The value of d_t was selected to be impulsive. 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 with a standard deviation of 0.00001. 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. □

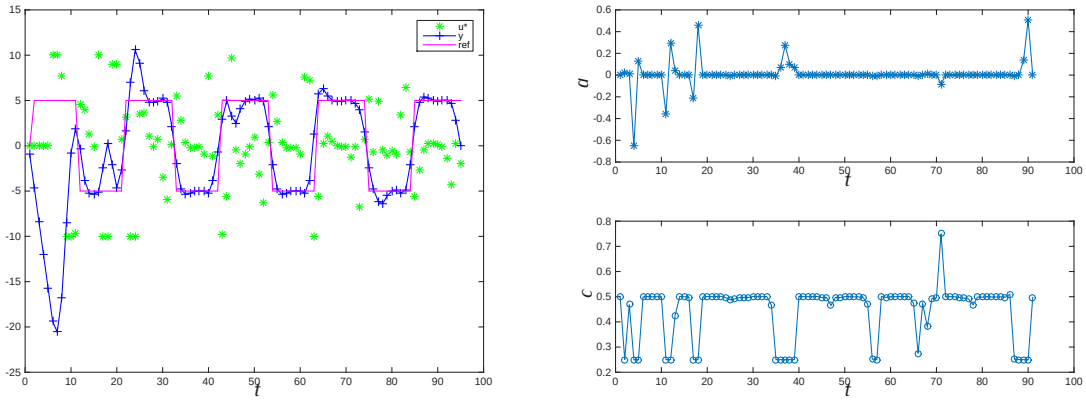

(a) Results of model uncertainty example. The measured output is in blue, the control sequence in green, and the reference signal in red.

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

Figure 6: Results of model uncertainty example.

Example 3 (Model uncertainty)**.** In this example we consider a linear system with parameter uncertainty in both the system dynamics and the output equation. Consider a linear discrete time-invariant system given by

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

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 any value of a in this interval, the system is unstable.

For this example the control objective is for the output y_s to track a reference r_s, so we consider solving an optimization of the form

$$\min_{\hat{u}_{t:t+T-1} \in \mathcal{U}} \quad \max_{\hat{x}_{t-L} \in \mathcal{X}, \hat{a} \in \mathcal{A}, \hat{c} \in \mathcal{C}} \sum_{s=t}^{t+T} \|y_s - r_s\|_2^2 + \lambda_u \sum_{s=t}^{t+T-1} \|\hat{u}_s\|_2^2 - \lambda_n \sum_{s=t-L}^{t} \|n_s\|_2^2, \tag{16}$$

where $\mathcal{A} \coloneqq \{a \in \mathbb{R} : -1 \leqslant a \leqslant 1\}$ and $\mathcal{C} \coloneqq \{c \in \mathbb{R} : 0.25 \leqslant c \leqslant 0.75\}$. We include the uncertain parameters a and c as maximization variables in order to optimize over their worst-case values.

Figure 6 shows the results for this example solving the optimization (16) 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 (5) 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 with a standard deviation of 0.005. Figure 6 shows that the controller is able to successfully regulate the system to the reference signal despite uncertain knowledge of the system. □

Example 4 (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

$$\begin{aligned} x1_{t+1} &= x1_t + v\cos(\theta_t), \\ x2_{t+1} &= x2_t + v\sin(\theta_t), \\ \theta_{t+1} &= \theta_t + u_t, \end{aligned} \qquad \begin{aligned} z1_{t+1} &= z1_t + d1_t, \\ z2_{t+1} &= z2_t + d2_t, \end{aligned}$$

where v is a constant scalar corresponding to the velocity of the pursuer, $x_t = \begin{bmatrix} x1_t & x2_t \end{bmatrix}' \in \mathbb{R}^2$ 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, $z_t = \begin{bmatrix} z1_t & z2_t \end{bmatrix}' \in \mathbb{R}^2$ is the position of the evader at time t, and $d_t = \begin{bmatrix} d1_t & d2_t \end{bmatrix}' \in \mathbb{R}^2$ is the evader's "acceleration" at time t. We assume that the control input u_t is bounded by the positive constant u_{max}, and "acceleration" of the evader d_t is bounded by the positive constant d_{max}. The output of the system is given by $y_t = \begin{bmatrix} x_t & z_t \end{bmatrix}' + \begin{bmatrix} n1_t & n2_t \end{bmatrix}'$, where $n_t = \begin{bmatrix} n1_t & n2_t \end{bmatrix}' \in \mathbb{R}^2$ 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.

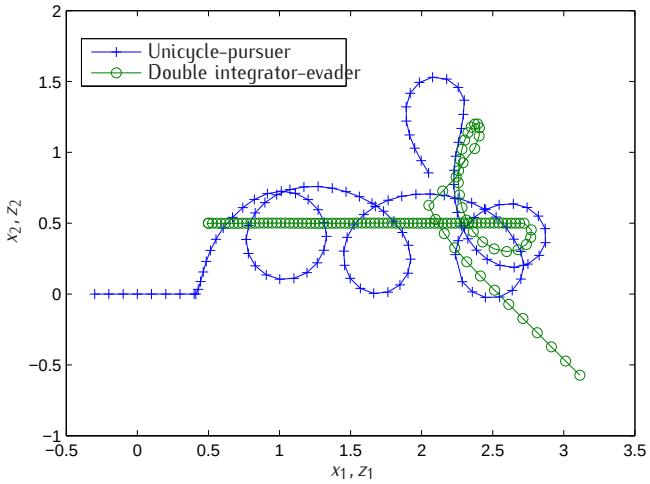

Figure 7: Results of nonlinear pursuit-evasion example. The path of the pursuer is shown with blue +'s, and the path of the evader is shown with green o's. Both players start on the left side of the figure.

Figure 7 shows the results for this problem solving the optimization (4) with the cost function given in (13) with $h_s(\cdot) = z_s - x_s$, $\mathcal{U} := \{u_t \in \mathbb{R} : -u_{max} \leqslant u_t \leqslant u_{max}\}$, $\mathcal{X} := \mathbb{R}^2 \times [0, 2\pi]$, 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 (5) 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^* 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^*. 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 while

Table 1: Numerical Performance for Examples 1-4

	Example 1	Example 2	Example 3	Example 4
# of optimization variables	103	174	29	157
# of inequality constraints	30	100	14	104
# of equality constraints	80	120	20	100
Mean time to compute	3.2 ms	5.1 ms	3.1 ms	2.6 ms
Max. time to compute	8.3 ms	9.7 ms	16.2 ms	5.0 ms
Min. time to compute	2.8 ms	4.3 ms	1.2 ms	1.9 ms

the evader is moving at a constant velocity, the pursuer continually catches up and is forced to loop back around the evader because of its unicycle-like dynamics. Once the evader starts playing "optimally" after time $t = 75$, with the given choices of parameter values, the evader is able to get away from the pursuer after a couple of maneuvers and is moving towards the right and bottom of Figure 4. □

Table 1 shows numerical performance for each of the examples above. The number of optimization variables, inequality and equality constraints, and the average, maximum, and minimum times that it takes for the solver to compute a solution at each time step are given. These times were reported when running simulations using the C programming language on a laptop with a 2.8GHz Intel® Core™ i7 processor[3]. As seen in the table, each iteration of the solver takes only a few milliseconds, which is fast enough to compute solutions to the min-max optimization online.

6 Conclusions

We reviewed a recent 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 main assumption that a saddle-point solution exists (which presumes standard controllability and observability), and the additional assumptions of reversible dynamics and a terminal cost that is an ISS-control Lyapunov function with respect to the disturbance input, Theorem 1 gives bounds on the state of the system and the tracking error for reference tracking problems.

We validated this approach by showing simulation results for both constrained linear and nonlinear examples. These examples included high-dimensional, uncertain, and unstable systems. We also showed that using our primal-dual-like interior-point method, we were able to solve large optimization problems very efficiently on a laptop, usually in just a few milliseconds.

Future work may involve investigating other types of noise and disturbances and implementing this online optimization approach in embedded applications. Also, relaxation of the saddle-point assumption to only require that an epsilon-close Nash equilibrium exists could be investigated.

References

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

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

[3]While the micro-processor used has multiple cores, a single core was used for the optimization.

[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] 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.

[15] D. A. Copp and J. P. Hespanha, "Nonlinear output-feedback model predictive control with moving horizon estimation," in *53rd IEEE Conference on Decision and Control*, pp. 3511–3517, IEEE, 2014.

[16] D. A. Copp and J. P. Hespanha, "Nonlinear output-feedback model predictive control with moving horizon estimation: Technical report," tech. rep., Univ. California, Santa Barbara, 2014.

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

[18] 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.

[19] 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.

[20] 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.

[21] 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.

[22] 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.

[23] 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.

[24] 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.

[25] 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.

[26] 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.

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

[28] 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.

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

[30] 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.

[31] 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.

[32] 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.

[33] 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.

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

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

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

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

[38] 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.

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

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

[41] 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.

[42] 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.

[43] 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.

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

[45] 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.

[46] 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.

[47] 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.

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

[49] 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.

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

[51] 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.