Automatica 48 (2012) 77–88

Contents lists available at SciVerse ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Stochastic receding horizon control with output feedback and bounded controls[☆]

Peter Hokayem[a,1], Eugenio Cinquemani[b], Debasish Chatterjee[c], Federico Ramponi[d], John Lygeros[a]

[a] Automatic Control Laboratory, ETH Zürich, Physikstrasse 3, 8092 Zürich, Switzerland

[b] INRIA Grenoble-Rhône-Alpes, 655 avenue de l'Europe, Montbonnot, 38 334 Saint Ismier Cedex, France

[c] Systems & Control Engineering, IIT Bombay, Powai, Mumbai 400076, India

[d] Department of Information Engineering, Università degli Studi di Brescia, via Branze 38, 25123 Brescia, Italy

ARTICLE INFO

Article history:
Received 1 October 2009
Received in revised form
17 March 2011
Accepted 12 July 2011
Available online 21 October 2011

Keywords:
Predictive control
Output feedback
Constrained control
State estimation
Stochastic control

ABSTRACT

We study the problem of receding horizon control for stochastic discrete-time systems with bounded control inputs and incomplete state information. Given a suitable choice of causal control policies, we first present a slight extension of the Kalman filter to estimate the state optimally in mean-square sense. We then show how to augment the underlying optimization problem with a negative drift-like constraint, yielding a second-order cone program to be solved periodically online. We prove that the receding horizon implementation of the resulting control policies renders the state of the overall system mean-square bounded under mild assumptions. We also discuss how some quantities required by the finite-horizon optimization problem can be computed off-line, thus reducing the on-line computation.

1. Introduction

A considerable amount of research has been devoted to deterministic receding horizon control, see, for example, (Maciejowski, 2001; Mayne, Rawlings, Rao, & Scokaert, 2000) and references therein. This resulted in proofs of recursive feasibility and stability of receding horizon control laws in the noise-free deterministic setting. These techniques can be extended to the robust case, i.e., whenever there is exogenous noise or parametric uncertainty of bounded nature entering the system. The counterpart for stochastic systems subject to process noise, imperfect state measurements, and bounded control inputs, however, is still lacking. The principal obstacle is posed by the fact that it may not be possible to determine an a priori bound on the support of the noise, for example, whenever the noise is additive and Gaussian. This extra ingredient complicates both the stability and the feasibility proofs: the noise, at least in the additive case, eventually drives the state outside any bounded set no matter how large the latter is taken to be, and employing any standard linear state feedback means that any hard bounds on the control inputs will eventually be violated.

In this article we propose a solution to the general receding horizon control problem for linear systems with noisy process dynamics, imperfect state information, and bounded control inputs. Both the process and measurement noise sequences are assumed to enter the system in an additive fashion, and we require that the designed control policies satisfy hard bounds. Periodically at times $t = 0, N_c, 2N_c, \ldots$, where N_c is the control horizon, a certain finite-horizon optimal control problem is solved over a prediction (or optimization) horizon $N \geqslant N_c$. The cost to be minimized is the standard expectation of the sum of cost-per-stage functions that are quadratic in the state and control inputs (Bertsekas, 2000, 2007). We can also include at the design level some variance-like bounds on the predicted future states and inputs—this is one possible way to impose soft state constraints that are in spirit similar to integrated chance-constraints, e.g., in Klein Haneveld (1983) and Klein Haneveld and van der Vlerk (2006).

There are several key challenges inherent to our setup. First, since the state information is imperfect one needs a filter to estimate the state. Second, in the presence of unbounded (e.g., Gaussian) noise, it is not possible in general to ensure any bound on the control values generated via linear state feedback; the additive nature of the noise ensures that the state exits from any fixed bounded set at some time almost surely, implying the necessity of nonlinear feedback policies. This issue is further complicated by the fact that only incomplete state information is available. Third, it is unclear whether the application of the

☆ This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor Fabrizio Dabbene under the direction of Editor Roberto Tempo.

E-mail addresses: hokayem@control.ee.ethz.ch (P. Hokayem), eugenio.cinquemani@inria.fr (E. Cinquemani), chatterjee@sc.iitb.ac.in (D. Chatterjee), federico.ramponi@ing.unibs.it (F. Ramponi), lygeros@control.ee.ethz.ch (J. Lygeros).

[1] Tel.: +41 44 632 0403; fax: +41 44 632 1211.

0005-1098/$ – see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2011.09.048

bounded control policies stabilizes the system in any reasonable sense. For a deterministic discrete-time linear system, $x_{t+1} = Ax_t + Bu_t$, it is not possible to globally asymptotically stabilize the system, if the matrix A has unstable eigenvalues, see, for example, Yang, Sontag, and Sussmann (1997) and references therein. Moreover, in the presence of stochastic process noise the hope for achieving asymptotic stability is obviously not realistic. In this article we relax the notion of stability to mean-square boundedness of the state and impose the extra conditions that the system matrix A is Lyapunov (or neutrally) stable and that the pair (A, B) is stabilizable.

The main contributions of this paper can be summarized as follows: Given a suitable subclass of bounded causal feedback policies, we show how to augment the finite-horizon optimal control problem to be solved periodically every N_c steps with a *stability constraint* and that the resulting optimization problem can be approximated to a globally feasible second-order cone program (SOCP). Under the assumption that the process and measurement noise are Gaussian, (even though the bounded inputs requirement makes the problem inherently nonlinear and the process statistics are non-Gaussian,) it turns out that Kalman filtering techniques can indeed be utilized. We rely on a low-complexity algorithm (essentially similar to standard Kalman filtering) for updating the conditional density of the state, given the history of the previous outputs, and report tractable solutions for the *off-line* computation of the time-dependent variance and covariance matrices in the optimization program. Finally, we show that the recursive application of the resulting control policies renders the state of the overall system mean-square bounded. This article builds upon and generalizes the earlier stability results in Hokayem, Chatterjee, Ramponi, Chaloulos, and Lygeros (2010) and Ramponi, Chatterjee, Milias-Argeitis, Hokayem, and Lygeros (2010) that were derived for the perfect state information case. Also, the current results generalize those in Chatterjee, Hokayem, and Lygeros (in press) and Hokayem, Chatterjee, and Lygeros (2009) from the perfect state information case to the imperfect state information case. In particular, if we have full state information available, the control policy proposed in this article reduces to that proposed in Hokayem et al. (2009) and to a special case of the policies proposed in Chatterjee et al. (in press) where the vector space is spanned by a single function. However, none of the earlier results, including those in Hokayem, Cinquemani, Chatterjee, and Lygeros (2010), are able to deal with recursive feasibility and stability for the setup of this article. Moreover, the control policy structure in this article is different from that in Hokayem, Cinquemani et al. (2010).

Related work

The research on stochastic receding horizon control is broadly subdivided into two parallel lines: the first treats multiplicative noise that enters the state equations, and the second caters to additive noise. The case of multiplicative noise has been treated in Couchman, Cannon, and Kouvaritakis (2006), Cannon, Kouvaritakis, and Wu (2009a) and Primbs and Sung (2009). In Primbs and Sung (2009), the noise enters the state equation multiplicatively and mixed hard state-input constraints are relaxed into expectation constraints. Terminal constraints are imposed as well that render the overall MPC scheme stable under full state feedback. The authors in Couchman et al. (2006) treat the case of uncertain output measurement matrix (C) and solve the MPC problem under probabilistic constraints on the outputs and full state feedback. In Cannon et al. (2009a) the stochastic MPC problem is treated under full state feedback and multiplicative noise entering the state equation. The proposed scheme comprises a pre-stabilizing linear state feedback control part and an open-loop part. The pre-stabilizing feedback gain is computed off-line and only the open-loop part is optimized online. The results in Cannon, Kouvaritakis, and Wu (2009b) extend those in Cannon et al. (2009a) to the case of additive noise as well. However, both results (Cannon et al., 2009a,b) involve a pre-stabilizing state feedback controller and hence no hard input bounds can be imposed.

We focus in this article on the additive noise case. The approach proposed here stems from and generalizes the idea of affine parametrization of control policies for finite-horizon linear quadratic problems proposed in Ben-Tal, Boyd, and Nemirovski (2006); Ben-Tal, Goryashko, Guslitzer, and Nemirovski (2004), utilized within the robust MPC framework in Ben-Tal et al. (2006), Goulart, Kerrigan, and Maciejowski (2006) and Löfberg (2003) for full state feedback, and in van Hessem and Bosgra (2003) for output feedback with Gaussian state and measurement noise inputs. More recently, this affine approximation was utilized in Skaf and Boyd (2010) for both the robust deterministic and the stochastic setups in the absence of control bounds; and optimality of affine policies in the scalar deterministic case was reported in Bertsimas, Iancu, and Parrilo (2010). In Bertsimas and Brown (2007) the authors reformulate the stochastic programming problem as a deterministic one with bounded noise support and solve a robust optimization problem over a finite horizon, followed by estimating the performance when the noise can take unbounded values, i.e., when the noise is unbounded, but takes high values with low probability. A similar approach was utilized in Oldewurtel, Jones, and Morari (2008) as well. There are also other approaches, e.g., those employing randomized algorithms as in Batina (2004), Blackmore (2006) and Maciejowski, Lecchini, and Lygeros (2007). Results on obtaining lower bounds on the value functions of the stochastic optimization problem have been recently reported in Wang and Boyd (2009), and a novel stochastic MPC scheme based on the scenario approach has appeared in Bernardini and Bemporad (2009). Other works employing probabilistic constraints may be found in Li, Wendt, and Wozny (2002) and Schwarm and Nikolaou (1999). In Magni, Pala, and Scattolini (2009), an input-to-state stability approach is employed and stability is shown under full state feedback and bounded additive process noise. An MPC scheme for systems with imperfect state information has been proposed in Yan and Bitmead (2005) under general hypotheses with probabilistic constraints. However, the ability to deal with noise of an unbounded nature (for example Gaussian) is still absent, in which stability and recursive feasibility could not be proven in Yan and Bitmead (2005) under bounded control inputs.

The rest of this article is organized as follows. We formulate the stochastic receding horizon control problem with all the underlying assumptions, the construction of the estimator, and the main optimization problem to be solved in Section 2. We provide the main results pertaining to tractability of the optimization problem and mean-square boundedness of the closed-loop system in Section 3. We comment on the obtained results and provide some extensions in Section 4. We then present numerical examples in Section 5 and conclude in Section 6. Finally, we provide the proofs in the Appendix.

Notation. Let $(\Omega, \mathfrak{F}, \mathbb{P})$ be a general probability space. We denote the conditional expectation given the sub-σ algebra $\mathfrak{F}'$ of $\mathfrak{F}$ as $\mathbb{E}_{\mathfrak{F}'}[.]$. For any random vector s we let $\Sigma_s := \mathbb{E}[ss^{\mathsf{T}}]$ and $\mathbb{E}_{\mathcal{Y}_t}[.]$ denote the conditional expectation given $\mathcal{Y}_t$. Hereafter we let $\mathbb{N}_+ := \{1, 2, \ldots\}$ and $\mathbb{N} := \mathbb{N}_+ \cup \{0\}$. We let $\operatorname{tr}(\cdot)$ denote the trace of a square matrix, $\|\cdot\|_p$ denote the standard p-norm, and simply $\|\cdot\|$ denote the Euclidean norm. We denote by $\|s\|_M := \sqrt{s^{\mathsf{T}} M s}$ for $M = M^{\mathsf{T}} \geqslant 0$. In a Euclidean space we denote by $\mathcal{B}_r$ the closed Euclidean ball of radius r centered at the origin. For any two matrices A and B of compatible dimensions, we denote by $\mathfrak{R}_k(A, B)$ the k-th step reachability matrix $\mathfrak{R}_k(A, B) := \left[A^{k-1}B \ \cdots \ AB \ B\right]$.

For any matrix M, we let $\sigma_{\min}$ and $\sigma_{\max}$ be its minimal and maximal singular values, respectively. We let $(M)_{i_1:i_2}$ denote the sub-matrix obtained by selecting the rows i_1 through i_2 of M and $(M)_i$ denote the i-th row of M. For any positive real number s, $\lceil s\rceil$ denotes the smallest integer that upper-bounds s.

2. Problem setup

Consider the following affine discrete-time stochastic dynamical model:

$$x_{t+1} = Ax_t + Bu_t + w_t, \tag{1a}$$

$$y_t = Cx_t + v_t, \tag{1b}$$

where $t \in \mathbb{N}$, $x_t \in \mathbb{R}^n$ is the state, $u_t \in \mathbb{R}^m$ is the control input, $y_t \in \mathbb{R}^p$ is the output, $w_t \in \mathbb{R}^n$ is a random process noise, $v_t \in \mathbb{R}^p$ is a random measurement noise, and A, B, and C are known matrices. We posit the following standing assumption:

Assumption 1. (i) The pair (A, B) is stabilizable (Bernstein, 2009, Chapter 12).

(ii) The matrix A is Lyapunov stable (Bernstein, 2009, Chapter 12), i.e., the eigenvalues $\{\lambda_i(A) \mid i = 1, \ldots, n\}$ lie in the closed unit disc, and those eigenvalues $\lambda_j(A)$ with $|\lambda_j(A)| = 1$ have equal algebraic and geometric multiplicities.

(iii) The initial condition and the process and measurement noise vectors are mutually independent and normally distributed, i.e., $x_0 \sim \mathcal{N}(0, \Sigma_{x_0})$, $w_t \sim \mathcal{N}(0, \Sigma_w)$, and $v_t \sim \mathcal{N}(0, \Sigma_v)$, with $\Sigma_w > 0$ and $\Sigma_v > 0$.

(iv) $\left(A, \Sigma_w^{1/2}\right)$ is controllable and (A, C) is observable.

(v) The control inputs satisfy

$$\|u_t\|_\infty \leqslant U_{\max} \quad \forall t \in \mathbb{N}. \tag{2}$$

Without loss of generality, we assume that A is given in real Jordan canonical form. Indeed, given a linear system described by the system matrices $(\tilde{A}, \tilde{B})$, there exists a coordinate transformation in the state-space that brings the pair $(\tilde{A}, \tilde{B})$ to the pair (A, B), where A is in real Jordan form (Horn & Johnson, 1990, p. 150). In particular, choosing a suitable ordering of the Jordan blocks, we can ensure that the pair (A, B) has the form $\left(\begin{bmatrix} A_s & 0 \\ 0 & A_o\end{bmatrix}, \begin{bmatrix} B_s \\ B_o\end{bmatrix}\right)$, where $A_s \in \mathbb{R}^{n_s\times n_s}$ is Schur stable, and $A_o \in \mathbb{R}^{n_o\times n_o}$ has its eigenvalues on the unit circle. By Assumption 1(ii), A_o is therefore block-diagonal with the diagonal blocks being either ± 1, or 2×2 rotation matrices. As a consequence, A_o is orthogonal. Moreover, since (A, B) is stabilizable, the pair (A_o, B_o) must be reachable in a number of steps $\kappa \leqslant n_o$ that depends on the dimension of A_o. Therefore, we can start by considering that the state Eq. (1a) has the form

$$\begin{bmatrix} x_{t+1}^s \\ x_{t+1}^o \end{bmatrix} = \begin{bmatrix} A_s x_t^s \\ A_o x_t^o \end{bmatrix} + \begin{bmatrix} B_s \\ B_o \end{bmatrix} u_t + \begin{bmatrix} w_t^s \\ w_t^o \end{bmatrix}, \tag{3}$$

where A_s is Schur stable, A_o is orthogonal, and there exists a nonnegative integer $\kappa \leqslant n_o$ such that the subsystem (A_o, B_o) is reachable in κ steps. This reachability index κ is fixed throughout the rest of the article.

For each $t \in \mathbb{N}$, let $\mathcal{Y}_t := \{y_0, \ldots, y_t\}$ denote the set of output observations up to time t. Fix a prediction horizon $N \in \mathbb{N}_+$, with $N \geqslant \kappa$, and define the cost $\mathcal{J}_t$ as

$$\mathcal{J}_t = \mathbb{E}_{\mathcal{Y}_t}\left[\sum_{k=0}^{N-1}\left(\|x_{t+k}\|^2_{Q_k} + \|u_{t+k}\|^2_{R_k}\right) + \|x_{t+N}\|^2_{Q_N}\right], \tag{4}$$

where $Q_k = Q_k^\mathsf{T} \geqslant 0$, $Q_N = {Q_N}^\mathsf{T} \geqslant 0$, and $R_k = R_k^\mathsf{T} \geqslant 0$ are given matrices of appropriate dimension, for $k = 0, \ldots, N-1$.

The evolution of the system (1a)–(1b) over a single prediction horizon N, starting at t, can be described in a compact form as

$$X_t = \mathcal{A}x_t + \mathcal{B}U_t + \mathcal{D}W_t, \qquad Y_t = \mathcal{C}X_t + V_t, \tag{5}$$

where $X_t = \begin{bmatrix} x_t \\ x_{t+1} \\ \vdots \\ x_{t+N}\end{bmatrix}$, $U_t = \begin{bmatrix} u_t \\ u_{t+1} \\ \vdots \\ u_{t+N-1}\end{bmatrix}$, $W_t = \begin{bmatrix} w_t \\ w_{t+1} \\ \vdots \\ w_{t+N-1}\end{bmatrix}$, $Y_t = \begin{bmatrix} y_t \\ y_{t+1} \\ \vdots \\ y_{t+N}\end{bmatrix}$, $V_t = \begin{bmatrix} v_t \\ v_{t+1} \\ \vdots \\ v_{t+N}\end{bmatrix}$, $\mathcal{A} = \begin{bmatrix} I \\ A \\ \vdots \\ A^N\end{bmatrix}$, $\mathcal{B} = \begin{bmatrix} 0 & \cdots & \cdots & 0 \\ B & \ddots & & \vdots \\ AB & B & \ddots & \vdots \\ \vdots & & \ddots & 0 \\ A^{N-1}B & \cdots & AB & B\end{bmatrix}$, $\mathcal{D} = \begin{bmatrix} 0 & \cdots & \cdots & 0 \\ I & \ddots & & \vdots \\ A & I & \ddots & \vdots \\ \vdots & & \ddots & 0 \\ A^{N-1} & \cdots & A & I\end{bmatrix}$, and $\mathcal{C} = \operatorname{diag}\{C, \ldots, C\}$. The cost function (4) at time t can also be written compactly as

$$\mathcal{J}_t = \mathbb{E}_{\mathcal{Y}_t}\left[\|X_t\|^2_{\mathcal{Q}} + \|U_t\|^2_{\mathcal{R}}\right], \tag{6}$$

where $\mathcal{Q} = \operatorname{diag}\{Q_0, \ldots, Q_N\}$ and $\mathcal{R} = \operatorname{diag}\{R_0, \ldots, R_{N-1}\}$. The cost $\mathcal{J}_t$ in (4) is a conditional expectation given the observations up to time t, the evaluation of which requires the conditional density $f(x_t \mid \mathcal{Y}_t)$ of the state given the previous and current measurements. For $t, s \in \mathbb{N}$, define $\hat{x}_{t|s} = \mathbb{E}_{\mathcal{Y}_s}[x_t]$ and $P_{t|s} = \mathbb{E}_{\mathcal{Y}_s}[(x_t - \hat{x}_{t|s})(x_t - \hat{x}_{t|s})^\mathsf{T}]$.

The following result is a slight extension of the standard Kalman filter. A proof may be found in Kumar and Varaiya (1986, p.102).

Proposition 2. *Let Assumption 1(iii) hold and assume that u_t is a deterministic function of $\mathcal{Y}_t$. Then $f(x_t \mid \mathcal{Y}_t)$ and $f(x_{t+1} \mid \mathcal{Y}_t)$ are the probability densities of Gaussian distributions $\mathcal{N}(\hat{x}_{t|t}, P_{t|t})$ and $\mathcal{N}(\hat{x}_{t+1|t}, P_{t+1|t})$, respectively, with $P_{t|t} \geqslant 0$ and $P_{t+1|t} \geqslant 0$. For $t = -1, 0, 1, 2, \ldots$, their conditional means and covariances can be computed iteratively starting at $(\hat{x}_{0|-1}, P_{0|-1}) := (0, \Sigma_{x_0})$, as follows:*

$$\begin{aligned}\hat{x}_{t+1|t+1} &= \hat{x}_{t+1|t} + P_{t+1|t}C^\mathsf{T}(CP_{t+1|t}C^\mathsf{T} \\ &\quad + \Sigma_v)^{-1}(y_{t+1} - C\hat{x}_{t+1|t})\end{aligned} \tag{7}$$

$$P_{t+1|t+1} = P_{t+1|t} - P_{t+1|t}C^\mathsf{T}(CP_{t+1|t}C^\mathsf{T} + \Sigma_v)^{-1}CP_{t+1|t} \tag{8}$$

$$\hat{x}_{t+1|t} = A\hat{x}_{t|t} + Bu_t \tag{9}$$

$$P_{t+1|t} = AP_{t|t}A^\mathsf{T} + \Sigma_w. \tag{10}$$

Proposition 2 states that the conditional mean and covariances of x_t can be propagated by an iterative algorithm which resembles the Kalman filter. In particular, the matrix $P_{t|t}$ together with $\hat{x}_{t|t}$ characterize the conditional density $f(x_t \mid \mathcal{Y}_t)$, which is needed in the computation of the cost (4) (or equivalently (6)). We note here that in the receding horizon control case considered in this paper u_t will be a nonlinear function of $\{y_0, \ldots, y_t\}$; therefore we cannot assume that all the probability distributions in the problem are Gaussian as in the case of LQG; in fact, the a priori distributions of x_t and of $\mathcal{Y}_t$ are not. Hereafter, we shall denote for notational convenience by $\hat{x}_t$ the estimate $\hat{x}_{t|t}$, and let

$$\hat{x}_t = \left[(\hat{x}_t^s)^\mathsf{T} \quad (\hat{x}_t^o)^\mathsf{T}\right]^\mathsf{T}, \tag{11}$$

which corresponds to the Jordan decomposition in (3). Let $K_t := (AP_{t|t}A^\mathsf{T} + \Sigma_w)C^\mathsf{T}(C(AP_{t|t}A^\mathsf{T} + \Sigma_w)C^\mathsf{T} + \Sigma_v)^{-1}$ and define $\Gamma_t := I - K_tC$ and $\Phi_t := \Gamma_tA$. Then, we can write the estimation error vector over a single prediction horizon N as

$$E_t := X_t - \hat{X}_t = \mathcal{F}_t e_t + \mathcal{G}_t W_t - \mathcal{H}_t V_t, \tag{12}$$

where $e_t = x_t - \hat{x}_t$, $\hat{X}_t = \begin{bmatrix} \hat{x}_t \\ \vdots \\ \hat{x}_{t+N} \end{bmatrix}$, $\mathcal{F}_t = \begin{bmatrix} I \\ \Phi_t \\ \Phi_{t+1} \cdot \Phi_t \\ \vdots \\ \Phi_{t+N-1} \cdots \Phi_t \end{bmatrix}$,

$$\mathcal{G}_t = \begin{bmatrix} 0 & \cdots & 0 & 0 \\ \Gamma_t & \cdots & 0 & 0 \\ \Phi_{t+1}\Gamma_t & \cdots & 0 & 0 \\ \vdots & & \vdots & \vdots \\ \Phi_{t+N-2} \cdots \Phi_{t+1}\Gamma_t & \cdots & \Gamma_{t+N-2} & 0 \\ \Phi_{t+N-1} \cdots \Phi_{t+1}\Gamma_t & \cdots & \Phi_{t+N-1}\Gamma_{t+N-2} & \Gamma_{t+N-1} \end{bmatrix},$$

$$\mathcal{H}_t = \begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ 0 & K_t & \cdots & 0 & 0 \\ 0 & \Phi_{t+1}K_t & \cdots & 0 & 0 \\ \vdots & \vdots & & \vdots & \vdots \\ 0 & \Phi_{t+N-2} \cdots \Phi_{t+1}K_t & \cdots & K_{t+N-2} & 0 \\ 0 & \Phi_{t+N-1} \cdots \Phi_{t+1}K_t & \cdots & \Phi_{t+N-1}K_{t+N-2} & K_{t+N-1} \end{bmatrix}.$$

The innovation sequence can be written as

$$Y_t - \hat{Y}_t = \mathcal{C}\mathcal{F}_t e_t + \mathcal{C}\mathcal{G}_t W_t + (I - \mathcal{C}\mathcal{H}_t)V_t, \tag{13}$$

where $\hat{Y}_t := \mathcal{C}\hat{X}_t$. Consequently, *the innovation sequence over the prediction horizon is independent of the input vector U_t*. Also, under the forgoing assumptions, the error vector e_t is a Gaussian random variable with mean zero and variance $P_{t|t}$.

2.1. Optimization problem and control policies

We would like to minimize the cost (4) over the class of all causal feedback policies. However, this optimization problem is extremely difficult to solve in general (Bertsekas, 2000, 2007). Therefore, we restrict attention to a subclass of causal feedback policies for which the optimization problem is tractable. Guided by our earlier approach in Chatterjee et al. (in press), Hokayem et al. (2009), Hokayem, Chatterjee et al. (2010), and Hokayem, Cinquemani et al. (2010) and given a control horizon $N_c \geqslant 1$ and a prediction horizon $N \geqslant N_c$, we would like to periodically minimize the cost (4) at times $t = 0, N_c, 2N_c, \ldots$ over the following class of control policies

$$u_{t+\ell} = \eta_{t+\ell} + \sum_{i=0}^{\ell} \theta_{t+\ell,t+i}\varphi_i(y_{t+i} - \hat{y}_{t+i}), \tag{14}$$

where $\ell = 0, 1, \ldots, N-1$, $\hat{y}_i = C\hat{x}_i$ is the output of the estimator, and for any vector $z = \big(z_1, \ldots, z_p\big) \in \mathbb{R}^p$, $\varphi_i(z) = \big(\varphi_{i,1}(z_1), \ldots, \varphi_{i,p}(z_p)\big)$, where $\varphi_{i,j} : \mathbb{R} \to \mathbb{R}$ is any function with $\sup_{s\in\mathbb{R}} |\varphi_{i,j}(s)| \leqslant \varphi_{\max} < \infty$ for some $\varphi_{\max} > 0$. The feedback gains $\theta_{\ell,i} \in \mathbb{R}^{m\times p}$ and the affine terms $\eta_\ell \in \mathbb{R}^m$ are the decision variables. The value of $u_{t+\ell}$ in (14) depends on the values of the measured outputs from the beginning of the prediction horizon at time t up to time $t + \ell$ only, which requires *finite memory*. Note that we have chosen to saturate the measurements we obtain from the vectors $(y_i - \hat{y}_i)$ before employing them in the control policy. This allows us to consider unbounded noise and yet ensure bounded policies; neither the process noise nor the measurement noise distributions are defined over a compact domain, in contrast to robust deterministic receding horizon control (Mayne et al., 2000) or other stochastic receding horizon control approaches as in Cannon, Cheng, Kouvaritakis, and Raković (2010). Moreover, the choice of element-wise saturation functions $\varphi_i(\cdot)$ is left open. As such, we can accommodate standard saturation, piecewise linear, and sigmoidal functions, to name a few. The control policy (14) at time t can be compactly written as

$$U_t = \boldsymbol{\eta}_t + \boldsymbol{\Theta}_t\varphi(Y_t - \hat{Y}_t), \tag{15}$$

where $\boldsymbol{\Theta}_t$ has the following (causal) lower block triangular structure

$$\boldsymbol{\Theta}_t := \begin{bmatrix} \theta_{t,t} & 0 & \cdots & 0 \\ \theta_{t+1,t} & \theta_{t+1,t+1} & & \vdots \\ \vdots & \vdots & \ddots & 0 \\ \theta_{t+N-1,t} & \theta_{t+N-1,t+1} & \cdots & \theta_{t+N-1,t+N-1} \end{bmatrix}, \tag{16}$$

$$\boldsymbol{\eta}_t := \begin{bmatrix} \eta_t \\ \eta_{t+1} \\ \vdots \\ \eta_{t+N-1} \end{bmatrix},$$

and $\varphi(Y_t - \hat{Y}_t) := \begin{bmatrix} \varphi_0(y_t - \hat{y}_t) \\ \vdots \\ \varphi_{N-1}(y_{t+N-1} - \hat{y}_{t+N-1}) \end{bmatrix}$.

Since the innovation vector $Y_t - \hat{Y}_t$ in (13) is not a function of $\boldsymbol{\eta}_t$ and $\boldsymbol{\Theta}_t$, the control inputs U_t in (14) remain affine in the decision variables. This fact is important to show convexity of the optimization problem, as will be seen in the next section. Finally, the constraint (2) can be rewritten as:

$$\|U_t\|_\infty \leqslant U_{\max} \quad \forall t = 0, N_c, 2N_c, \ldots. \tag{17}$$

Summarizing, the optimization problem to be solved periodically at times $t = 0, N_c, 2N_c, \ldots$ is given by

$$\min_{(\boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)} \Big\{ \mathcal{J}_t \;\big|\; (5), (15), (16), (17) \Big\}. \tag{18}$$

3. Main results

Even if problem (18) is successively feasible every N_c steps, in general the resulting control actions do not guarantee stability of the resulting receding horizon controller. Unlike standard deterministic stability arguments utilized in MPC, see, for example, (Mayne et al., 2000), we *cannot* assume the existence of a compact robust positively invariant terminal region, since the process noise sequence does not have a compact support. Instead, we introduce an additional *stability constraint* which, if recursively feasible, renders the state of the closed-loop system mean-square bounded. Guided by the argument in Ramponi et al. (2010), we then show that this constraint is indeed recursively feasible.

For $t = 0, N_c, 2N_c, \ldots$, the state estimate at time $t + N_c$ can be written as

$$\hat{x}_{t+N_c} = A^{N_c}\hat{x}_t + \mathfrak{R}_{N_c}(A, B)\begin{bmatrix} u_t \\ \vdots \\ u_{t+N_c-1} \end{bmatrix} + \Xi_t, \tag{19}$$

where $\mathfrak{R}_{N_c}(A, B)$ is the reachability matrix as defined earlier and Ξ_t is defined as

$$\begin{aligned} \Xi_t := {} & \begin{bmatrix} A^{N_c-1}K_tCA & A^{N_c-2}K_{t+1}CA & \cdots & K_{t+N_c-1}CA \end{bmatrix}\begin{bmatrix} e_t \\ \vdots \\ e_{t+N_c-1} \end{bmatrix} \\ & + \begin{bmatrix} A^{N_c-1}K_tC & A^{N_c-2}K_{t+1}C & \cdots & K_{t+N_c-1}C \end{bmatrix}\begin{bmatrix} w_t \\ \vdots \\ w_{t+N_c-1} \end{bmatrix} \\ & + \begin{bmatrix} A^{N_c-1}K_t & A^{N_c-2}K_{t+1} & \cdots & K_{t+N_c-1} \end{bmatrix}\begin{bmatrix} v_{t+1} \\ \vdots \\ v_{t+N_c} \end{bmatrix}. \end{aligned} \tag{20}$$

In order to show boundedness of the state variance, we require that the N_c-step iteration (19) has bounded variance. However, this estimate has the term Ξ_t, which involves the error in the state estimation process as well as the process and measurement noise vectors. This term Ξ_t may be viewed as 'noise' entering the system, with bounded fourth moment. In particular, we require that there exists, at least after some time T', a uniform bound on its first moment. This is captured by the following Proposition.

Proposition 3. *There exists an integer T' and a positive constant ζ, depending on the given problem parameters, such that*

$$\mathbb{E}_{\mathcal{Y}_t}\left[\|\Xi_t\|\right] \leqslant \zeta \quad \text{for all } t \geqslant T'. \tag{21}$$

Using the constant ζ, we now require the following "drift condition" to be satisfied: for any chosen constant $\varepsilon > 0$ and for every $t = 0, N_c, 2N_c, \ldots, U_t \in \mathbb{U}$ is designed such that the following condition is satisfied

$$\left\| A_o^{N_c}\hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)\begin{bmatrix} u_t \\ \vdots \\ u_{t+N_c-1}\end{bmatrix}\right\| \leqslant \left\|\hat{x}_t^o\right\| - \left(\zeta + \frac{\varepsilon}{2}\right)$$
$$\text{whenever } \left\|\hat{x}_t^o\right\| \geqslant \zeta + \varepsilon. \tag{22}$$

As will be shown later, condition (22) above guarantees that on average the state norm *contracts* every N_c steps, a crucial ingredient towards showing mean-square boundedness of the closed-loop system. Moreover, N_c needs to be chosen appropriately (depending on the reachability index κ) in order to ensure that the constraint is feasible. Note that $\begin{bmatrix} u_t \\ \vdots \\ u_{t+N_c-1}\end{bmatrix} = (\boldsymbol{\eta}_t)_{1:N_c m} + (\boldsymbol{\Theta}_t)_{1:N_c m}\varphi(Y - \hat{Y})$. (For notational convenience, we have retained $\varphi(Y - \hat{Y})$ with the knowledge that the matrix $(\boldsymbol{\Theta}_t)_{1:N_c m}$ causally selects the first N_c output vectors as they become available, see (16).) We augment problem (18) with the stability constraint (22) to obtain

$$\min_{(\boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)} \left\{ \mathcal{J}_t \;\middle|\; (5), (15), (16), (17), (22)\right\}. \tag{23}$$

The ingredients of our stochastic receding horizon control problem corresponding to (23) are summarized in Algorithm 1.

Algorithm 1 Basic Stochastic Receding Horizon Algorithm

Require: density $f(x_0|\mathcal{Y}_{-1}) := \mathcal{N}(0, \Sigma_{x_0})$
set $t \leftarrow 0, \hat{x}_{0|-1} \leftarrow 0$, and $P_{0|-1} \leftarrow \Sigma_{x_0}$
loop
  for $i = 0$ to $N_c - 1$ **do**
    measure y_{t+i}
    calculate $\hat{x}_{t+i}(= \hat{x}_{t+i|t+i})$ and $P_{t+i|t+i}$ using (7)–(8)
    if $i = 0$ **then**
      solve the optimization problem (23) for the optimal policy $\{u_t^*, \ldots, u_{t+N-1}^*\}$
    end if
    using the obtained control policy above, compute and apply u_{t+i}^*
    calculate $\hat{x}_{t+i+1|t+i}$ and $P_{t+i+1|t+i}$ using (9)–(10)
  end for
  set $t \leftarrow t + N_c$
end loop

Assumption 4. We require that:

(i) The control and prediction horizons satisfy $N \geqslant N_c = \kappa$, where κ is the reachability index of the orthogonal subsystem (A_o, B_o) in (3).

(ii) The control authority $U_{\max} \geqslant U_{\max}^*$, where $U_{\max}^* := \sigma_{\min}(\mathfrak{R}_{N_c}(A_o, B_o))^{-1}\left(\zeta + \frac{\varepsilon}{2}\right)$ and $\mathfrak{R}_{N_c}(A_o, B_o)$ is the N_c-step reachability matrix of the orthogonal subsystem.

In fact, choosing any control horizon $N_c \geqslant \kappa$ turns out to be sufficient in order to have a feasible control vector U_t for problem (23) with an upper bound $U_{\max} = U_{\max}^*$; however, we will take $N_c = \kappa$ for simplicity.

Theorem 5. *Consider the system (1a)–(1b), and suppose that Assumptions 1 and 4 hold. Then:*

(i) *For every time $t = 0, N_c, 2N_c, \ldots,$ the optimization problem (23) in Algorithm 1 is convex and can be conservatively approximated and solved via the following globally (hence recursively) feasible second-order cone program (SOCP):*

$$\underset{(z_1, z_2, z_3, \boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)}{\text{minimize}} \quad z_1 \tag{24}$$

subject to

$$\left\|\boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \Lambda_t^{\varphi}\right\|_{\mathcal{M}}^2 + \operatorname{tr}\left(\boldsymbol{\Theta}_t^{\mathsf{T}}\mathcal{M}\boldsymbol{\Theta}_t(\Lambda_t^{\varphi\varphi} - \Lambda_t^{\varphi}\Lambda_t^{\varphi\mathsf{T}})\right) + 2\hat{x}_t^{\mathsf{T}}\mathcal{A}^{\mathsf{T}}\mathcal{Q}\mathcal{B}\boldsymbol{\eta}_t + 2\operatorname{tr}\left(\boldsymbol{\Theta}_t^{\mathsf{T}}\mathcal{B}^{\mathsf{T}}\mathcal{Q}(\mathcal{D}\Lambda_t^{w\varphi} + \mathcal{A}\Lambda_t^{x\varphi})\right) \leqslant z_1 \tag{25}$$

$$|(\boldsymbol{\eta}_t)_i| + \|(\boldsymbol{\Theta}_t)_i\|_1 \varphi_{\max} \leqslant U_{\max} \quad \forall i = 1, \ldots, Nm, \tag{26}$$

$$\left.\begin{aligned} &\left\|A_o^{N_c}\hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\eta}_t)_{1:N_c m}\right\| \leqslant z_2 \\ &\left\|\mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\Theta}_t)_{1:N_c m}\right\|_{\infty} \leqslant z_3 \\ &z_2 + \sqrt{n_o}\varphi_{\max} z_3 \leqslant \left\|\hat{x}_t^o\right\| - \left(\zeta + \frac{\varepsilon}{2}\right) \end{aligned}\right\}$$
$$\text{whenever } \left\|\hat{x}_t^o\right\| \geqslant \zeta + \varepsilon \tag{27}$$

the structure of $\boldsymbol{\Theta}_t$ in (16), where $\mathcal{M} := \mathcal{R} + \mathcal{B}^{\mathsf{T}}\mathcal{Q}\mathcal{B}$, and

$$\begin{aligned} &\Lambda_t^{\varphi} := \mathbb{E}_{\mathcal{Y}_t}[\varphi(Y_t - \hat{Y}_t)], \qquad \Lambda_t^{x\varphi} := \mathbb{E}_{\mathcal{Y}_t}[x_t\varphi(Y_t - \hat{Y}_t)^{\mathsf{T}}], \\ &\Lambda_t^{w\varphi} := \mathbb{E}_{\mathcal{Y}_t}[W\varphi(Y_t - \hat{Y}_t)^{\mathsf{T}}], \\ &\Lambda_t^{\varphi\varphi} := \mathbb{E}_{\mathcal{Y}_t}[\varphi(Y_t - \hat{Y}_t)\varphi(Y_t - \hat{Y}_t)^{\mathsf{T}}]. \end{aligned} \tag{28}$$

(ii) *The application of Algorithm 1 via the SOCP approximation in part (i) above renders the closed-loop system mean-square bounded, i.e., for any initial $\mathcal{Y}_0$, there exists a (computable) finite constant $\gamma > 0$, depending on the given problem parameters, such that*

$$\sup_{t \in \mathbb{N}} \mathbb{E}_{\mathcal{Y}_0}\left[\|x_t\|^2\right] \leqslant \gamma. \tag{29}$$

In practice, it may be also of interest to further impose constraints both on the state and the input vectors. For example, one may be interested in imposing linear and/or quadratic constraints on the state of the form

$$\mathbb{E}_{\mathcal{Y}_t}\left[\|X_t\|_{\mathcal{S}}^2 + \mathcal{L}^{\mathsf{T}}X_t\right] \leqslant \alpha_t, \tag{30}$$

where $\mathcal{S} = \mathcal{S}^{\mathsf{T}} \geqslant 0$ and $\alpha_t > 0$. Moreover, expected energy expenditure constraints can be posed as follows

$$\mathbb{E}_{\mathcal{Y}_t}\left[\|U_t\|_{\tilde{\mathcal{S}}}^2\right] \leqslant \beta_t, \tag{31}$$

where $\tilde{\mathcal{S}} = \tilde{\mathcal{S}}^{\mathsf{T}} \geqslant 0$ and $\beta_t > 0$. In the absence of hard input constraints, such expectation-type constraints are commonly used in the stochastic MPC (Eugenio, Agarwal, Chatterjee, D & Lygeros, in press; Primbs & Sung, 2009) and in stochastic optimization in the form of integrated chance constraints (Klein Haneveld, 1983; Klein Haneveld & van der Vlerk, 2006). This is partly because it is not possible, without posing further restrictions on the boundedness of the process noise w_t, to ensure that hard constraints on the state are satisfied. For example, in the standard LQG setting nontrivial hard constraints on the system state would generally be violated with nonzero probability. Moreover, in contrast to chance constraints where a bound is imposed on the probability of constraint violation, expectation-type constraints tend to give rise to convex optimization problems under weak assumptions (Eugenio et al., in press; Klein Haneveld, 1983; Klein Haneveld & van der Vlerk, 2006). We can augment problem (23) with the constraints (30) and (31) to obtain

$$\min_{(\eta_t,\Theta_t)} \left\{ \mathcal{J}_t \mid (5), (15), (16), (17), (22), (30), (31) \right\}. \tag{32}$$

Notice that the constraints (30) and (31) are not necessarily feasible at time t for any choice of parameters α_t and β_t. As such, problem (32) may become infeasible over time if we simply apply Algorithm 1. We therefore modify step 7 in Algorithm 1 to generate Algorithm 2. In this new version of the algorithm, problem (32) is either feasible at step 7 with the given α_t and β_t, or a bisection search is implemented (steps 12–26), with α_t and β_t as lower bounds, and upper bounds

$$\begin{aligned}\alpha_t^* &:= 3\,\mathrm{tr}\left(\mathcal{A}^\mathsf{T}\mathcal{S}\mathcal{A}\mathbb{E}_{\mathcal{Y}_t}\left[x_t {x_t}^\mathsf{T}\right] + \mathcal{D}^\mathsf{T}\mathcal{S}\mathcal{D}\Sigma_w\right) + \mathcal{L}^\mathsf{T}\mathcal{A}\hat{x}_t \\ &\quad + 3Nm\sigma_{\max}(\mathcal{B}^\mathsf{T}\mathcal{S}\mathcal{B})U_{\max}^2 + \left\|\mathcal{L}^\mathsf{T}\mathcal{B}\right\|_1 U_{\max}\end{aligned}$$

$$\beta_t^* := Nm\sigma_{\max}(\tilde{\mathcal{S}})U_{\max}^2 \tag{33}$$

that guarantee feasibility. The search is iterated until the change in α and β falls below a pre-specified precision number δ or a maximal number of iterations $\bar{\nu}$ is reached, which is used to keep the computational burden limited.

Corollary 6. *Consider the system (1a)–(1b), and suppose that Assumptions 1 and 4 hold. Then:*

(i) *For every time $t = 0, N_c, 2N_c, \ldots$ the optimization problem (32) in Algorithm 2 is convex and can be conservatively approximated and solved via the following globally (hence recursively) feasible (at either step 7 or step 13) second-order cone program (SOCP):*

$$\underset{(z_1,z_2,z_3,\eta_t,\Theta_t)}{\text{minimize}} \quad z_1$$

subject to

$$\begin{aligned}&\left\|\eta_t + \Theta_t\Lambda_t^\varphi\right\|_{\mathcal{M}}^2 + \mathrm{tr}\left({\Theta_t}^\mathsf{T}\mathcal{M}\Theta_t(\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi\Lambda_t^{\varphi\mathsf{T}})\right) \\ &\quad + 2{\hat{x}_t}^\mathsf{T}\mathcal{A}^\mathsf{T}\mathcal{Q}\mathcal{B}\eta_t + 2\,\mathrm{tr}\left({\Theta_t}^\mathsf{T}\mathcal{B}^\mathsf{T}\mathcal{Q}(\mathcal{D}\Lambda_t^{w\varphi} + \mathcal{A}\Lambda_t^{x\varphi})\right) \leqslant z_1 \\ &|(\eta_t)_i| + \|(\Theta_t)_i\|_1\,\varphi_{\max} \leqslant U_{\max} \quad \forall i = 1,\ldots,Nm, \\ &\left.\begin{aligned}&\left\|A_o^{N_c}\hat{x}_t^o + \mathfrak{R}_{N_c}(A_o,B_o)(\eta_t)_{1:N_c m}\right\| \leqslant z_2 \\ &\left\|\mathfrak{R}_{N_c}(A_o,B_o)(\Theta_t)_{1:N_c m}\right\|_\infty \leqslant z_3 \\ &z_2 + \sqrt{n_o}\varphi_{\max}z_3 \leqslant \left\|\hat{x}_t^o\right\| - \left(\zeta + \frac{\varepsilon}{2}\right)\end{aligned}\right\} \\ &\qquad \textit{whenever } \left\|\hat{x}_t^o\right\| \geqslant \zeta + \varepsilon \\ &\left\|\eta_t + \Theta_t\Lambda_t^\varphi\right\|_{\mathcal{B}^\mathsf{T}\mathcal{S}\mathcal{B}}^2 + \mathrm{tr}\left({\Theta_t}^\mathsf{T}\mathcal{B}^\mathsf{T}\mathcal{S}\mathcal{B}\Theta_t(\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi\Lambda_t^{\varphi\mathsf{T}})\right) \\ &\quad + 2{\hat{x}_t}^\mathsf{T}\mathcal{A}^\mathsf{T}\mathcal{S}\mathcal{B}\eta_t + 2\,\mathrm{tr}\left({\Theta_t}^\mathsf{T}\mathcal{B}^\mathsf{T}\mathcal{S}(\mathcal{D}\Lambda_t^{w\varphi} + \mathcal{A}\Lambda_t^{x\varphi})\right) \\ &\quad + \mathcal{L}^\mathsf{T}\mathcal{B}(\eta_t + \Theta_t\Lambda_t^\varphi) + \mathrm{tr}\left(\mathcal{A}^\mathsf{T}\mathcal{S}\mathcal{A}\mathbb{E}_{\mathcal{Y}_t}\left[x_t{x_t}^\mathsf{T}\right]\right) \\ &\quad + \mathrm{tr}\left(\mathcal{D}^\mathsf{T}\mathcal{S}\mathcal{D}\Sigma_w\right) + \mathcal{L}^\mathsf{T}\mathcal{A}\hat{x}_t \leqslant \alpha_t \end{aligned} \tag{34}$$

$$\left\|\eta_t + \Theta_t\Lambda_t^\varphi\right\|_{\tilde{\mathcal{S}}}^2 + \mathrm{tr}\left({\Theta_t}^\mathsf{T}\tilde{\mathcal{S}}\Theta_t(\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi\Lambda_t^{\varphi\mathsf{T}})\right) \leqslant \beta_t \tag{35}$$

the structure of Θ_t in (16), where all the required constant matrices are defined as in Theorem 5.

(ii) *The application of Algorithm 2 via the SOCP approximation in part* (i) *above renders the closed-loop system mean-square bounded, i.e., for any initial $\mathcal{Y}_0$, there exists a (computable) finite constant $\gamma > 0$, depending on the given problem parameters, such that*

$$\sup_{t\in\mathbb{N}} \mathbb{E}_{\mathcal{Y}_0}\left[\|x_t\|^2\right] \leqslant \gamma. \tag{36}$$

Algorithm 2 Modified Stochastic Receding Horizon Algorithm

Require: density $f(x_0|\mathcal{Y}_{-1}) := \mathcal{N}(0, \Sigma_{x_0})$
1: set $t \leftarrow 0$, $\hat{x}_{0|-1} \leftarrow 0$, and $P_{0|-1} \leftarrow \Sigma_{x_0}$
2: **loop**
3: **for** $i = 0$ to $N_c - 1$ **do**
4: measure y_{t+i}
5: calculate $\hat{x}_{t+i} (= \hat{x}_{t+i|t+i})$ and $P_{t+i|t+i}$ using (7)–(8)
6: **if** $i = 0$ **then**
7: solve the optimization problem (32) using the given α_t and β_t
8: **if** step 7 is feasible **then**
9: save the optimal sequence $\{u_t^*, u_{t+1}^*, \ldots, u_{t+N-1}^*\}$
10: goto step 28
11: **else**
12: set $\overline{\alpha} \leftarrow \alpha_t^*$, $\underline{\alpha} \leftarrow \alpha_t$, $\overline{\beta} \leftarrow \beta_t^*$, and $\underline{\beta} \leftarrow \beta_t$
13: solve the optimization problem (32) using $\overline{\alpha}$ and $\overline{\beta}$ to obtain the sequence $\{u_t^*, u_{t+1}^*, \ldots, u_{t+N-1}^*\}$
14: set $\nu \leftarrow 1$
15: **repeat**
16: set $\alpha_t \leftarrow (\overline{\alpha} + \underline{\alpha})/2$ and $\beta_t \leftarrow (\overline{\beta} + \underline{\beta})/2$
17: solve the optimization problem (32) using the new α_t and β_t
18: **if** step 17 is feasible **then**
19: set $\overline{\alpha} \leftarrow \alpha_t$ and $\overline{\beta} \leftarrow \beta_t$
20: save the new optimal sequence $\{u_t^*, u_{t+1}^*, \ldots, u_{t+N-1}^*\}$
21: **else**
22: set $\underline{\alpha} \leftarrow \alpha_t$ and $\underline{\beta} \leftarrow \beta_t$
23: **end if**
24: set $\nu = \nu + 1$
25: **until** $(|\overline{\alpha} - \underline{\alpha}| \le \delta$ and $|\overline{\beta} - \underline{\beta}| \le \delta)$ or $(\nu > \bar{\nu})$
26: **end if**
27: **end if**
28: apply u_{t+i}^*
29: calculate $\hat{x}_{t+i+1|t+i}$ and $P_{t+i+1|t+i}$ using (9)–(10)
30: **end for**
31: set $t = t + N_c$
32: **end loop**

4. Discussion

4.1. Recursive feasibility

The SOCPs solved in Theorem 5 and Corollary 6 are globally feasible, independently of the initial conditions of the plant and the estimator. As such, there is no a priori requirement for an initially feasible and invariant set of initial conditions, as is the case in nominal or robust receding horizon control (Mayne et al., 2000). This guarantee of recursive feasibility is shown in the proofs of Theorem 5 and Corollary 6 by providing a feasible control law that satisfies all the constraints in the SOCPs.

4.2. Mean-square boundedness

The mean-square boundedness conditions (29) and (36) provide an on-average guarantee that the state does not grow arbitrarily large. This is a weaker notion of stability than, for example, asymptotic stability or input-to-state stability (ISS) that have been utilized in nominal and robust receding horizon control, respectively. However, in the presence of possibly unbounded process and measurement noise, it is virtually impossible to guarantee that the state converges to the origin or that it is ultimately bounded in some compact set for every initial condition and every realization of the noise processes. In this case, similarly to LQG (i.e., even with unbounded control authority), mean-square boundedness is the best that can be aimed for within our setting, given the unboundedness of the noise processes and the limited control authority. The constants γ in (29) and (36) may be computed using the derivations in the Appendix and the formulas in the remark in Pemantle and Rosenthal (1999, pp.145).

4.3. More general policies

It is not difficult to show that one can also use quadratic policies of the form

$$U_t = \boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \varphi(Y_t - \hat{Y}_t) + \tilde{\boldsymbol{\Theta}}_t \tilde{\varphi}(Y_t - \hat{Y}_t), \quad (37)$$

instead of (14), where $\tilde{\boldsymbol{\Theta}}_t$ has the same causal structure of of $\boldsymbol{\Theta}_t$ and $\tilde{\varphi}(Y_t - \hat{Y}_t) := [\tilde{\varphi}_0(y_t - \hat{y}_t)^\mathsf{T}\tilde{\varphi}_0(y_t - \hat{y}_t) \cdots \tilde{\varphi}_{N-1}(y_{t+N-1} - \hat{y}_{t+N-1})^\mathsf{T} \tilde{\varphi}_{N-1}(y_{t+N-1} - \hat{y}_{t+N-1})]^\mathsf{T}$, with $\sup_{s\in\mathbb{R}} \mid \tilde{\varphi}_i(s) \mid \leqslant \tilde{\varphi}_{\max} < \infty$ for some $\tilde{\varphi}_{\max} > 0$. The underlying optimization problems (23) and (32) with the policy (37) are still convex and both Theorem 5 and Corollary 6 still apply with minor changes.

4.4. Off-line computation of the Λ matrices

The optimization problems (23) and (32) solved in Theorem 5 and in Corollary 6, respectively, are second-order cone programs (SOCP) for which efficient numerical solvers are available via software packages such as `yalmip` (Löfberg, 2004). As such, the optimization may be performed online. However, at any time $t = 0, N_c, 2N_c, \ldots$, our ability to solve the optimization problems in Theorem 5 and Corollary 6, respectively, hinges upon the computation of the matrices in (28).

Recall that $Y_t - \hat{Y}_t$ is the innovation sequence that was given in (13), and that $\hat{x}_t$ is the optimal mean-square estimate of x_t given the history $\mathcal{Y}_t$. The matrices (28) may be computed by numerical integration with respect to the independent Gaussian measures of $w_t, \ldots, w_{t+N-1}$, of $v_t, \ldots, v_{t+N}$, and of $(x_t - \hat{x}_t)$ given $\mathcal{Y}_t$. Due to the large dimensionality of the integration space, this approach may be impractical for online computations. One alternative approach relies on the observation that Λ_t^{φ}, $\Lambda_t^{w\varphi}$, and $\Lambda_t^{\varphi\varphi}$ depend on x_t via the difference $x_t - \hat{x}_t$. Since $x_t - \hat{x}_t$ is conditionally zero-mean given $\mathcal{Y}_t$, we can write the dependency of (28) on the time-varying statistics of x_t given $\mathcal{Y}_t$ as follows: $\Lambda_t^{x\varphi}(\hat{x}_t, P_{t|t}) = \Lambda_t^{e\varphi}(P_{t|t}) + \hat{x}_t \Lambda_t^{\varphi}(P_{t|t})^\mathsf{T}$, $\Lambda_t^{w\varphi}(P_{t|t})$, and $\Lambda_t^{\varphi\varphi}(P_{t|t})$, where $\Lambda_t^{e\varphi} := \mathbb{E}_{\mathcal{Y}_t}[(x_t - \hat{x}_t)\varphi(Y_t - \hat{Y}_t)^\mathsf{T}]$. In principle one may compute *off-line* and store the matrices $\Lambda_t^{e\varphi}(P_{t|t})$, $\Lambda_t^{\varphi}(P_{t|t})$, $\Lambda_t^{w\varphi}(P_{t|t})$, and $\Lambda_t^{\varphi\varphi}(P_{t|t})$, which depend on the covariance matrices $P_{t|t}$ but not on $\hat{x}_t$, and just update online the value of $\Lambda_t^{x\varphi}(\hat{x}_t, P_{t|t})$ as the estimate $\hat{x}_t$ becomes available. However, this poses serious requirements in terms of memory. A more appealing alternative is to exploit the convergence properties of the covariance matrix $P_{t|t}$. The following result can be inferred, for instance, from Kamen and Su (1999, Theorem 5.1).

Proposition 7. *Under Assumption 1*(iii) *and* (iv) *the discrete-time algebraic Riccati equation in* $P \in \mathbb{R}^{n\times n}$, $P = A[P - PC^\mathsf{T}(CPC^\mathsf{T} + \Sigma_v)^{-1}CP]A^\mathsf{T} + \Sigma_w$, *has a unique solution* $P^* \geqslant 0$. *The sequence* $P_{t+1|t}$ *defined by* (8) *and* (10) *converges to* P^* *as* t *tends to* ∞, *for any initial condition* $P_{0|-1} \geqslant 0$.

As a consequence, from (8) one sees that $P_{t|t}$ converges to $P^\circ = P^* - P^*C^\mathsf{T}(CP^*C^\mathsf{T} + \Sigma_v)^{-1}CP^*$, which is the asymptotic error covariance matrix of the estimator $\hat{x}_t$. Thus, neglecting the initial transient, one may just compute off-line and store the matrices $\Lambda_t^{e\varphi}(P^\circ)$, $\Lambda_t^{\varphi}(P^\circ)$, $\Lambda_t^{w\varphi}(P^\circ)$, and $\Lambda_t^{\varphi\varphi}(P^\circ)$, and just update the matrix $\Lambda_t^{x\varphi}(\hat{x}_t, P^\circ)$ for new values of the estimate $\hat{x}_t$.

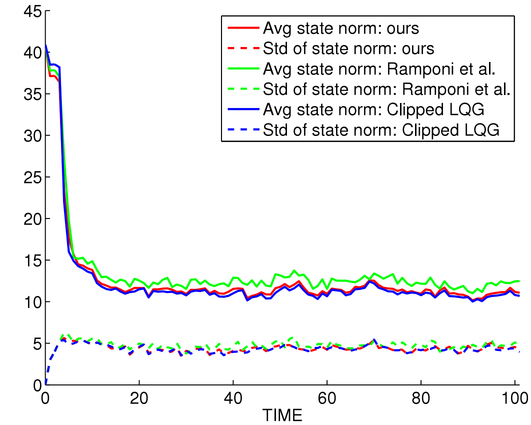

Fig. 1. Average and standard deviation of the state norm for $U_{\max} = 453$.

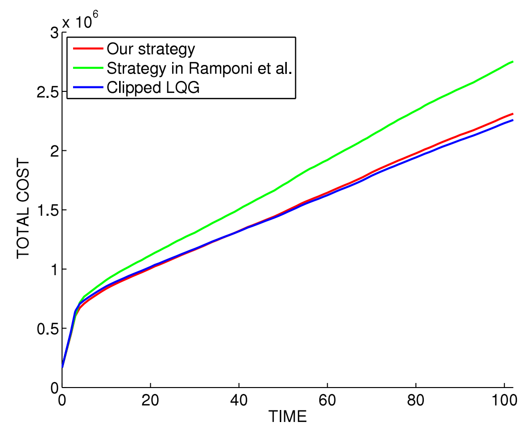

Fig. 2. Total cost for $U_{\max} = 453$.

5. Simulations

Consider the system (1a)–(1b) with the following matrices: $A = \begin{bmatrix} 0.9 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos(\psi) & -\sin(\psi) \\ 0 & 0 & \sin(\psi) & \cos(\psi) \end{bmatrix}$, $B = \begin{bmatrix} 0 \\ 1 \\ 0 \\ 1 \end{bmatrix}$, and $C = I$, where $\psi = \frac{\pi}{2}$. The orthogonal part of the state is 3-dimensional, and the controllability index of the orthogonal part is $\kappa = 3$.

Example 1. The simulation data was chosen to be: $x_0 \sim \mathcal{N}(0, I)$, $w_t \sim \mathcal{N}(0, 10I)$, $v_t = \mathcal{N}(0, 10I)$, $Q = I$, $R = 1$, $N = 5$, $N_c = \kappa = 3$, and φ the usual piecewise linear saturation function with $\varphi_{\max} = 1$. For this example the theoretical bound on the input is $U_{\max} \simeq 453$ for a choice of $\varepsilon = 10$.

We simulated the system above using Algorithm 1 for the discrete-time interval [0, 100]. In comparison, we simulated also the policy (38) proposed by the authors in Ramponi et al. (2010) whenever the state is estimated using a Kalman filter, with the understanding that this goes beyond the results in Ramponi et al. (2010), since Ramponi et al. (2010) deals only with the case of perfect state information. We also simulated the standard LQG controller for this system with post-saturation of the obtained controls. The average state norm as well as the standard deviation of the state norm using the three strategies are depicted in Fig. 1 and the total costs are plotted in Fig. 2. Fig. 2 shows approximately 16% improvement in the cost after 100 time steps by using Algorithm 1 versus the policy (38) in Ramponi et al. (2010) coupled with a Kalman filter. The performance of our policy is close to

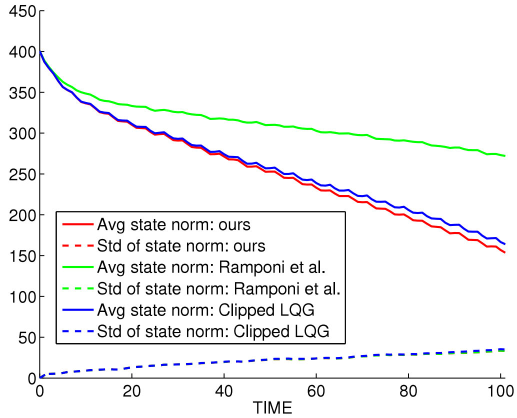

Fig. 3. Average and standard deviation of the state norm for $U_{\max} = 3.2664$.

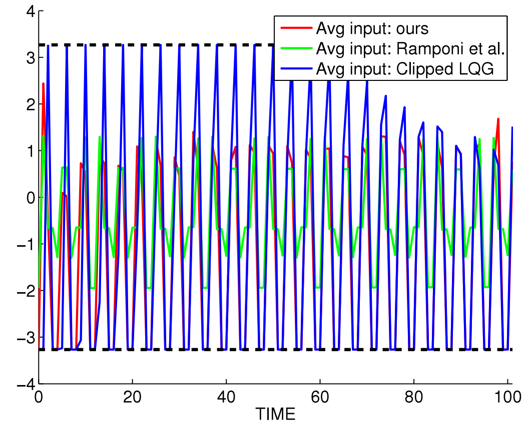

Fig. 5. Average inputs for $U_{\max} = 3.2664$.

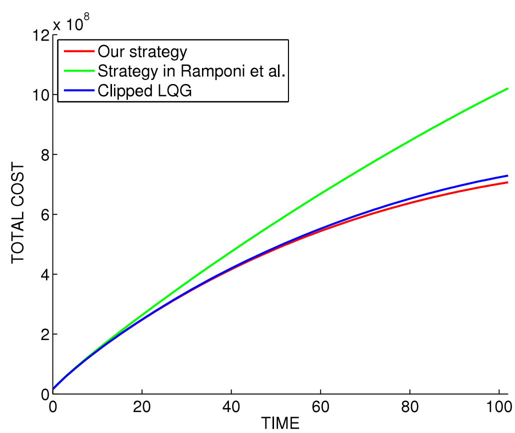

Fig. 4. Total cost for $U_{\max} = 3.2664$.

that of clipped LQG, which remains mean-square 'optimal' since the theoretical upper bound on the control authority is never reached by the LQG policy, i.e., the LQG policy is never clipped (for our choice of initial condition x_0). The optimization problem was solved using `yalmip` (Löfberg, 2004) and `sdpt3` (Toh, Todd, & Tatuncu, 1999) and the solver times were as follows: minimum = 0.1238 s, maximum = 0.5831 s, average = 0.1691 s, and standard deviation = 0.0144 s. This data pertains to an Intel(R) Core(TM)2 Duo CPU running at 2.66 GHz, with a Linux operating system. The computation of the matrices $\Lambda^{e\varphi}(P^\circ)$, $\Lambda^{\varphi}(P^\circ)$, $\Lambda^{w\varphi}(P^\circ)$, and $\Lambda^{\varphi\varphi}(P^\circ)$ was done off-line using the steady state error covariance matrix P°, as discussed in the previous section, via classical Monte Carlo integration using 10^5 samples.

Example 2. We also simulated the same system as before with $Q = 100I$ and $R = 1$. In this case the theoretical bound $\zeta = 336.6$ and the corresponding $U_{\max} = 440.45$. As this theoretical bound on ζ is conservative, we reduced ζ down to 2 and chose $\varepsilon = 0.5$, which result in $U_{\max} = 3.2664$. This choice of ζ is far below the required theoretical bound (which was given in Example 1 above) and as such the stability guarantees in the article do not apply anymore, i.e., there is no theoretical guarantee that the closed-loop system is mean-square bounded. However, it is apparent from Figs. 3 and 4 that our policy is stabilizing. It is important to notice that the clipped LQG policy hits the saturation level $U_{\max}$ quite often, whereas our policy does not, as seen in Fig. 5. However, despite the fact that our strategy does not take full advantage of the available control authority, it is still able to outperform the clipped LQG as well as the adapted policy in Ramponi et al. (2010).

6. Conclusions

We presented a method for stochastic receding horizon control of discrete-time linear systems with process and measurement noise and bounded input policies. We showed that the optimization problem solved periodically is successively feasible and convex. Moreover, we illustrated how a certain stability condition can be utilized to ensure that the application of the receding horizon controller renders the state of the system mean-square bounded. We discussed how certain matrices in the cost function can be computed off-line and provided examples that illustrate our approach, showing conditions under which it outperforms certain competing approaches.

Acknowledgments

This research was partially supported by the Swiss National Science Foundation under grant 200021-122072 and by the European Commission under the projects Feednetback FP7-ICT-223866 (www.feednetback.eu) and MoVeS FP7-ICT-2009-257005.

Appendix. Proofs

We begin by considering the estimation equation in (7) and combining it with (9) and the system dynamics (1a)–(1b) to obtain (recall that $\hat{x}_t$ is used instead of $\hat{x}_{t|t}$ for notational simplicity) $\hat{x}_{t+1} = A\hat{x}_t + Bu_t + K_t\big(CA(x_t - \hat{x}_t) + Cw_t + v_{t+1}\big)$, where $K_t = \left(AP_{t|t}A^\top + \Sigma_w\right)C^\top\left(C(AP_{t|t}A^\top + \Sigma_w)C^\top + \Sigma_v\right)^{-1}$ and $P_{t|t}$ is the error covariance matrix defined in (8). Both K_t and P_t are uniformly norm-bounded as shown in the following Lemma:

Lemma 8. *Consider the system (1a)–(1b), and let Assumption 1(iv) hold. In addition, assume that $P_{0|0} \geqslant 0$. Then, there exists a time $T' \in \mathbb{N}$ and constants $\rho, \rho_m > 0$ such that $\operatorname{tr}\left(P_{t|t}\right) \leqslant \rho$ and $\|K_t\| \leqslant \rho_m, \forall t \geqslant T'$.*

Proof. First, observe that $\sum_{i=0}^{\kappa_1-1} A^i\Sigma_w(A^i)^\top = \left[\Sigma_w^{1/2} \; A\Sigma_w^{1/2} \cdots A^{\kappa_1-1}\Sigma_w^{1/2}\right]\left[\Sigma_w^{1/2} \; A\Sigma_w^{1/2} \; \cdots \; A^{\kappa_1-1}\Sigma_w^{1/2}\right]^\top$, and since $(A, \Sigma_w^{1/2})$ is controllable by Assumption 4(iv), we see that there exists $\kappa_1 \in \mathbb{N}_+$ such that for all $k \geqslant \kappa_1$ the rank of $\left[\Sigma_w^{1/2} \; A\Sigma_w^{1/2} \; \cdots \; A^{\kappa_1-1}\Sigma_w^{1/2}\right] =$

n; indeed, κ_1 is the reachability index of $(A, \Sigma_w^{1/2})$. Thus, $\sum_{i=0}^{\kappa_1-1} A^i \Sigma_w (A^i)^\top$ is positive definite, and therefore, there exists some $\delta_1', \delta_2' > 0$ such that $\delta_1' I \leqslant \sum_{i=0}^{\kappa_1-1} A^i \Sigma_w (A^i)^\top \leqslant \delta_2' I$. Second, observe that $\sum_{i=0}^{\kappa_2-1} (A^i)^\top C^\top \Sigma_v^{-1} CA^i = \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{\kappa_2-1} \end{bmatrix}^\top \left(I_{\kappa_2} \otimes \Sigma_v^{-1}\right) \begin{bmatrix} C \\ CA \\ \vdots \\ CA^{\kappa_2-1} \end{bmatrix}$, where $\otimes$ denotes the standard Kronecker product. Since (A, C) is observable by assumption, there exists $\kappa_2 \in \mathbb{N}_+$ such that the rank of the matrix $\begin{bmatrix} C^\top & A^\top C^\top & \cdots & (A^{\kappa_2-1})^\top C^\top \end{bmatrix}^\top$ is n. The matrix $I_{\kappa_2} \otimes \Sigma_v^{-1}$ is clearly positive definite by Assumption 4(iv), and therefore, we see that there exists $\delta_1'', \delta_2'' > 0$ such that $\delta_1'' I \leqslant \sum_{i=0}^{\kappa_2-1} (A^i)^\top C^\top \Sigma_v^{-1} CA^i \leqslant \delta_2'' I$. Third, the conditions of Lemma 7.1 in Jazwinski (1970, pp. 234) are satisfied, for a choice of $\delta_1 = \min\{\delta_1', \delta_1''\}$ and $\delta_2 = \max\{\delta_2', \delta_2''\}$, and the bound $P_{t|t} \leqslant \rho' I$ for some $\rho' > 0$ is established for all $t \geqslant T' := \max\{\kappa_1, \kappa_2\}$. The assertion now follows immediately from: $\mathbb{E}_{\mathcal{Y}_t}\left[\left\|x_t - \hat{x}_t\right\|^2\right] = \operatorname{tr}\left(P_{t|t}\right) \leqslant n\lambda_{\max}(P_{t|t}) \leqslant n\rho' =: \rho$. Since by assumption $\Sigma_v > 0$, one possible bound on $\left(\|K_t\|\right)_{t \geqslant T'}$ is given by $\|K_t\| \leqslant \frac{\|AP_{t|t}A^\top + \Sigma_w\| \|C^\top\|}{\lambda_{\min}(\Sigma_v)} \leqslant \frac{\left(\|\Sigma_w\| + \|A\|^2 \|P_{t|t}\|\right)\|C^\top\|}{\lambda_{\min}(\Sigma_v)} =: \rho_m$. □

Using the bounds in Lemma 8, we can proceed to prove Proposition 3.

Proof (*Proof of* Proposition 3)**.** Recall the expression of Ξ_t in (20) and define the following quantities:
$F_t := \left[A^{N_c-1} K_t CA \quad A^{N_c-2} K_{t+1} CA \quad \cdots \quad K_{t+N_c-1} CA\right]$,
$G_t := [A^{N_c-1} K_t C \; A^{N_c-2} K_{t+1} C \; \cdots \; K_{t+N_c-1} C]$,
$H_t := [A^{N_c-1} K_t \; A^{N_c-2} K_{t+1} \; \cdots \; K_{t+N_c-1}]$.
Using Lemma 8, we have that $\|F_t\| \leqslant N_c \|CA\| \|K_t\| \leqslant N_c \rho_m \|CA\|$, $\|G_t\| \leqslant N_c \|C\| \|K_\ell\| \leqslant N_c \rho_m \|C\|$, and $\|H_t\| \leqslant N_c \|K_\ell\| \leqslant N_c \rho_m$, for all $t \geqslant T'$. It follows that

$$\begin{aligned} \mathbb{E}_{\mathcal{Y}_t}\left[\|\Xi_t\|\right] &\leqslant N_c \rho_m \|CA\| \, \mathbb{E}_{\mathcal{Y}_t}\left[\left\| \begin{matrix} e_t \\ \vdots \\ e_{t+N_c-1} \end{matrix} \right\|\right] \\ &\quad + N_c \rho_m \|C\| \, \mathbb{E}_{\mathcal{Y}_t}\left[\left\| \begin{matrix} w_t \\ \vdots \\ w_{t+N_c-1} \end{matrix} \right\|\right] + N_c \rho_m \mathbb{E}_{\mathcal{Y}_t}\left[\left\| \begin{matrix} v_{t+1} \\ \vdots \\ v_{t+N_c} \end{matrix} \right\|\right] \\ &\leqslant N_c{}^{3/2} \rho_m (\|CA\| \sqrt{\rho} + \|C\| \sqrt{\operatorname{tr}(\Sigma_w)} + \sqrt{\operatorname{tr}(\Sigma_v)}) =: \zeta, \\ &\quad \text{for all} t \geqslant T'. \quad \square \end{aligned}$$

Proof of Theorem 5. We begin by showing the first claim in Theorem 5. □

Lemma 9. *Consider the system* (1a)–(1b)*, and suppose that Assumption* 4 *holds. Then assertion*(i) *of* Theorem 5 *holds.*

Proof. *Convexity*: It is clear that $X_t{}^\top \mathcal{Q} X_t + U_t{}^\top \mathcal{R} U_t$ is convex in X_t and U_t, and both X_t and U_t are affine functions of the design parameters $(\boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)$ for every realization of the noise sequences $(w_t)_{t\in\mathbb{N}}$ and $(v_t)_{t\in\mathbb{N}}$. Since taking expectation of a convex function retains convexity (Boyd & Vandenberghe, 2004), we conclude that the cost $V_t = \mathbb{E}_{\mathcal{Y}_t}\left[X_t{}^\top \mathcal{Q} X_t + U_t{}^\top \mathcal{R} U_t\right]$ is convex in $(\boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)$. Similarly, the constraints (17) and (22) are convex in $(\boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)$ as they are a composition of convex and affine functions (Boyd & Vandenberghe, 2004).

SOCP formulation: Substituting the augmented dynamics (5) into the objective function (6), we have that $\mathcal{J}_t = \mathbb{E}_{\mathcal{Y}_t}[\|\mathcal{A}x_t + \mathcal{B}U_t + \mathcal{D}W_t\|_{\mathcal{Q}}^2 + \|U_t\|_{\mathcal{R}}^2] = \mathbb{E}_{\mathcal{Y}_t}\left[\|U_t\|_{\mathcal{M}}^2 + 2U_t{}^\top \mathcal{B}^\top \mathcal{Q}(\mathcal{A}x_t + \mathcal{D}W_t)\right] + \mathbb{E}_{\mathcal{Y}_t}\left[\|W_t\|_{\mathcal{D}^\top\mathcal{Q}\mathcal{D}}^2 + \|x_t\|_{\mathcal{A}^\top\mathcal{Q}\mathcal{A}}^2\right]$, where we have used the fact that the noise W_t is zero-mean (and $\mathcal{M} = \mathcal{R} + \mathcal{B}^\top \mathcal{Q} \mathcal{B}$). Note that the last term above does not depend on the decision variables so we shall henceforth drop it from the optimization. Now, substituting the policy (15) into the last equation and completing the square yields $\mathcal{J}_t \tilde{=} \mathbb{E}_{\mathcal{Y}_t}\left[(\boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \varphi(Y_t - \hat{Y}_t))^\top \mathcal{M}(\boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \varphi(Y_t - \hat{Y}_t)) + 2(\boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \varphi(Y_t - \hat{Y}_t))^\top \mathcal{B}^\top \mathcal{Q}(\mathcal{A}x_t + \mathcal{D}W_t)\right] \left\|\boldsymbol{\eta}_t + \boldsymbol{\Theta}_t \Lambda_t^\varphi\right\|_{\mathcal{M}}^2 + \operatorname{tr}\left(\boldsymbol{\Theta}_t{}^\top \mathcal{M} \boldsymbol{\Theta}_t (\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi \Lambda_t^{\varphi\top})\right) + 2\boldsymbol{\eta}_t{}^\top \mathcal{B}^\top \mathcal{Q} \mathcal{A} \hat{x}_t + 2 \operatorname{tr}(\boldsymbol{\Theta}_t{}^\top \mathcal{B}^\top \mathcal{Q} (\mathcal{A}\Lambda_t^{x\varphi} + \mathcal{D}\Lambda_t^{w\varphi}))$. Note that $\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi \Lambda_t^{\varphi\top} \geqslant 0$, so we can further write the quadratic term in $\boldsymbol{\Theta}_t$ in a symmetric fashion, i.e., since $\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi \Lambda_t^{\varphi\top}$ is positive semidefinite, there exists a matrix L so that $\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi \Lambda_t^{\varphi\top} = LL^\top$. It follows that $\operatorname{tr}\left(\boldsymbol{\Theta}_t{}^\top \mathcal{M} \boldsymbol{\Theta}_t (\Lambda_t^{\varphi\varphi} - \Lambda_t^\varphi \Lambda_t^{\varphi\top})\right) = \operatorname{tr}\left(\boldsymbol{\Theta}_t{}^\top \mathcal{M} \boldsymbol{\Theta}_t LL^\top\right) = \operatorname{tr}(L)^\top \boldsymbol{\Theta}_t{}^\top \mathcal{M} \boldsymbol{\Theta}_t L$. Using the epigraph formulation gives the cost (24) and constraint (25).

Concerning the constraint (26), we have shown in Chatterjee et al. (in press) and Hokayem et al. (2009) that combining the constraint $\|u_t\|_\infty \leqslant U_{\max}$ and the class of policies (15) is equivalent to the constraints $\left|(\boldsymbol{\eta}_t)_i\right| + \|(\boldsymbol{\Theta}_t)_i\|_1 \varphi_{\max} \leqslant U_{\max}$ for all $i = 1, \ldots, Nm$. Substituting (15) into the stability constraint (22), we obtain

$$\begin{aligned} &\|A_o^{N_c} \hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\eta}_t)_{1:N_c m} \\ &\quad + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\Theta}_t)_{1:N_c m} \varphi(Y_t - \hat{Y}_t)\| \\ &\leqslant \left\|A_o^{N_c} \hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\eta}_t)_{1:N_c m}\right\| \\ &\quad + \left\|\mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\Theta}_t)_{1:N_c m} \varphi(Y_t - \hat{Y}_t)\right\| \\ &\leqslant \left\|A_o^{N_c} \hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\eta}_t)_{1:N_c m}\right\| \\ &\quad + \sqrt{n_o} \left\|\mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\Theta}_t)_{1:N_c m} \varphi(Y_t - \hat{Y}_t)\right\|_\infty \\ &\leqslant \left\|A_o^{N_c} \hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\eta}_t)_{1:N_c m}\right\| \\ &\quad + \sqrt{n_o} \left\|\mathfrak{R}_{N_c}(A_o, B_o)(\boldsymbol{\Theta}_t)_{1:N_c m}\right\|_\infty \varphi_{\max}. \end{aligned}$$

Enforcing that the last term above is $\leqslant \left\|\hat{x}_t^o\right\| - \left(\zeta - \frac{\varepsilon}{2}\right)$, whenever $\left\|\hat{x}_t^o\right\| \geqslant \zeta + \varepsilon$, is equivalent to the constraint (27), where the decision variables are now $(z_2, z_3, \boldsymbol{\eta}_t, \boldsymbol{\Theta}_t)$. Moreover, if the constraint (27) is satisfied, then the stability constraint (22) is satisfied as well. As such, the optimization problem solved in Theorem 5(i) is a conservative approximation of (23) due to the fact that the constraint (27) is tighter than (22).

Feasibility: It remains to show that all the constraints are simultaneously feasible. Inspired by the work in Ramponi et al. (2010), we consider the candidate controller

$$\tilde{u}_{t,t+N_c-1} = \begin{bmatrix} \eta_t \\ \vdots \\ \eta_{t+N_c-1} \end{bmatrix} = -\mathfrak{R}_{N_c}(A_o, B_o)^\dagger \operatorname{sat}_r(A_o^{N_c} \hat{x}_t^o) \tag{38}$$

with $\eta_{t+N_c} = \eta_{t+N_c+1} = \cdots = \eta_{t+N-1} = 0$ and $\boldsymbol{\Theta}_t = 0$, where $r := \zeta + \varepsilon/2$ and recall that $\operatorname{sat}_r(v) = v$ if $\|v\| \leqslant r$ and $\operatorname{sat}_r(v) = rv/\|v\|$ if $\|v\| > r$. First, we have that $\left\|\tilde{u}_{t,t+N_c-1}\right\|_\infty \leqslant \left\|\tilde{u}_{t,t+N_c-1}\right\|_2 \leqslant \sigma_{\min}(\mathfrak{R}_{N_c}(A_o, B_o))^{-1}(\zeta + \varepsilon/2) = U_{\max}^*$, and the constraint (2) is feasible. Regarding the constraint (22), we have that $\left\|A_o^{N_c} \hat{x}_t^o + \mathfrak{R}_{N_c}(A_o, B_o)\tilde{u}_{t,t+N_c-1}\right\| = \left\|\hat{x}_t^o\right\| - r \leqslant \left\|\hat{x}_t^o\right\| - (\zeta + \varepsilon/2)$, whenever $\left\|\hat{x}_t^o\right\| \geqslant \zeta + \varepsilon$, where the first equality follows from the orthogonality of A_o (see Ramponi et al. (2010)), and the constraint (22) is also feasible.

Finally, it is easy to see that the cost (24) is linear and (25) through (27) are second-order cone constraints, hence the optimization program is a SOCP (Ben-Tal & Nemirovski, 2001). □

The following result pertains to mean-square boundedness of the Schur subsystem $\hat{x}^s_t$ of the estimator, i.e., $\hat{x}^s_{t+1} = A_s\hat{x}^s_t + B_s u_t + (K_t)_{1:n_s}\big(CA(x_t - \hat{x}_t) + Cw_t + v_{t+1}\big)$, where $(K_t)_{1:n_s}$ are the first n_s rows of the gain K_t.

Lemma 10. *Let Assumption 1 hold. Then there exists a constant $\gamma_s > 0$, depending on the given problem parameters, such that $\mathbb{E}_{\mathcal{Y}_{T'}}\big[\|\hat{x}^s_t\|^2\big] \leqslant \gamma_s, \forall t \geqslant T'$, where T' is as defined in Proposition 3.*

Proof. Since the matrix A_s is Schur stable, there exists a positive definite matrix $M \in \mathbb{R}^{n_s\times n_s}$ that satisfies $A_s{}^{\mathsf{T}}MA_s - M = -I$, see Bernstein (2009, Proposition 11.10.5). Pick a constant $\nu \in]0, \min\{1, 1/\lambda_{\max}(M)\}[$ such that $A_s{}^{\mathsf{T}}MA_s - M \leqslant -\nu M$. Then for any $t \in \mathbb{N}$, $\mathbb{E}_{\mathcal{Y}_t}\Big[\|\hat{x}^s_{t+1}\|^2_M\Big] - \|\hat{x}^s_t\|^2_M \leqslant -\nu\|\hat{x}^s_t\|^2_M + 2\mathbb{E}_{\mathcal{Y}_t}\big[(\hat{x}^s_t)^{\mathsf{T}}A_s{}^{\mathsf{T}}MB_s u_t\big] + \mathbb{E}_{\mathcal{Y}_t}[\|(K_t)_{1:n_s}\big(CA(x_t-\hat{x}_t)+Cw_t+v_{t+1}\big)\|^2_M]$, where $\|x\|^2_M := x^{\mathsf{T}}Mx$. Using Young's inequality and Assumption 1, we have that $2\mathbb{E}_{\mathcal{Y}_t}\big[(\hat{x}^s_t)^{\mathsf{T}}A_s{}^{\mathsf{T}}MB_s u_t\big] \leqslant \varepsilon\mathbb{E}_{\mathcal{Y}_t}\big[\|A_s\hat{x}^s_t\|^2_M\big] + \frac{1}{\varepsilon}\mathbb{E}_{\mathcal{Y}_t}\big[\|B_s u_t\|^2_M\big] \leqslant \varepsilon\mathbb{E}_{\mathcal{Y}_t}\big[\|A_s\hat{x}^s_t\|^2_M\big] + \frac{m\sigma_{\max}(B_s^TMB_s)}{\varepsilon}U^2_{\max}$, for $\varepsilon > 0$ and for all $t \in \mathbb{N}$. Also, using Lemma 8 and Assumption 1, we have that $\mathbb{E}_{\mathcal{Y}_t}\Big[\big\|(K_t)_{1:n_s}\big(CA(x_t-\hat{x}_t)+Cw_t+v_{t+1}\big)\big\|^2_M\Big] \leqslant 3\lambda_{\max}(M)\rho^2_m(\|CA\|^2\rho + \|C\|^2\,\mathrm{tr}(\Sigma_w) + \mathrm{tr}(\Sigma_v))$. Choose an $\varepsilon \leqslant \frac{\nu}{2\|A_s\|^2_M}$ and let $c := \frac{m\sigma_{\max}(B_s^TMB_s)}{\varepsilon}U^2_{\max} + 3\lambda_{\max}(M)\rho^2_m(\|CA\|^2\rho + \|C\|^2\,\mathrm{tr}(\Sigma_w) + \mathrm{tr}(\Sigma_v))$, then we have $\mathbb{E}_{\mathcal{Y}_t}\Big[\|\hat{x}^s_{t+1}\|^2_M\Big] \leqslant \big(1-\frac{\nu}{2}\big)\|\hat{x}^s_t\|^2_M + c$, for all $t \geqslant T'$. Iterating the last inequality, we have $\mathbb{E}_{\mathcal{Y}_{T'}}\Big[\|\hat{x}^s_t\|^2_M\Big] \leqslant \big(1-\frac{\nu}{2}\big)^{(t-T')}\|\hat{x}^s_{T'}\|^2_M + \sum_{i=0}^{t-T'-1}\big(1-\frac{\nu}{2}\big)^i c \leqslant \|\hat{x}^s_{T'}\|^2_M + \big(1-\frac{\nu}{2}\big)^{-1}c$, for all $t \geqslant T'$. Setting $\gamma_s := \|\hat{x}^s_{T'}\|^2_M + \big(1-\frac{\nu}{2}\big)^{-1}c$, completes the proof. □

We consider next the orthogonal subsystem of the estimator and show that the process $(\hat{x}^o_t)_{t\in\mathbb{N}}$ is mean-square bounded. We shall rely on the following fundamental result pertaining to mean-square boundedness of a general random sequence $(\xi_t)_{t\in\mathbb{N}}$; it is an immediate consequence of Pemantle and Rosenthal (1999, Theorem 1).

Proposition 11. *Let $(\xi_t)_{t\in\mathbb{N}}$ be a sequence of nonnegative random variables on some probability space $(\Omega, \mathfrak{F}, \mathbb{P})$, and let $(\mathfrak{F}_t)_{t\in\mathbb{N}}$ be any filtration to which $(\xi_t)_{t\in\mathbb{N}}$ is adapted. Suppose that there exist constants $\varepsilon > 0$, and $J, M < \infty$, such that $\xi_0 \leqslant J$, and for all $t \in \mathbb{N}$:*

$$\mathbb{E}_{\mathfrak{F}_t}\big[\xi_{t+1} - \xi_t\big] \leqslant -\frac{\varepsilon}{2} \quad \text{on the event } \{\xi_t > J\} \tag{39}$$

$$\text{and} \quad \mathbb{E}_{\xi_0,\ldots,\xi_t}\big[|\xi_{t+1} - \xi_t|^4\big] \leqslant M. \tag{40}$$

Then there exists a constant $\gamma = \gamma(\varepsilon, J, M) > 0$ such that $\sup_{t\in\mathbb{N}}\mathbb{E}\big[\xi_t^2\big] \leqslant \gamma$.

Lemma 12. *Let Assumptions 1 and 4 hold. Then there exists a constant $\gamma_o > 0$, depending on the given problem parameters, such that*

$$\mathbb{E}_{\mathcal{Y}_T}\Big[\|\hat{x}^o_t\|^2\Big] \leqslant \gamma_o, \quad \forall t \geqslant T := N_c\lceil T'/\kappa\rceil. \tag{41}$$

Proof. Consider the subsampled process $\hat{x}^o_{t+N_c} = A_o^{N_c}\hat{x}^o_t + \mathfrak{R}_{N_c}(A_o, B_o)u_{t,t+N_c-1} + (\Xi_t)_{n-n_o+1:n}$ for $t = 0, N_c, 2N_c, \ldots$, where Ξ_t is as defined in (20) and $u_{t,t+N_c-1} := \begin{bmatrix} u_t \\ \vdots \\ u_{t+N_c-1}\end{bmatrix}$. We shall first verify the two conditions (39) and (40) of Proposition 11 for the process $(\xi_t)_{t=0,N_c,2N_c,\ldots} = \big(\|\hat{x}^o_t\|\big)_{t=0,N_c,2N_c,\ldots}$. Using the triangle inequality, we have that $\mathbb{E}_{\mathcal{Y}_t}\Big[\|\hat{x}^o_{t+N_c}\| - \|\hat{x}^o_t\|\Big] \leqslant \mathbb{E}_{\mathcal{Y}_t}\Big[\big\|A_o^{N_c}\hat{x}^o_t + \mathfrak{R}_{N_c}(A_o,B_o)u_{t,t+N_c-1}\big\| - \|\hat{x}^o_t\|\Big] + \mathbb{E}_{\mathcal{Y}_t}\Big[\big\|(\Xi_t)_{n-n_o+1:n}\big\|\Big]$. We know from Proposition 3 that there exists a uniform (with respect to time t) upper bound ζ for the last term on the right-hand side of the last inequality for $t \geqslant T := N_c\lceil T'/\kappa\rceil$. Accordingly, we have $\mathbb{E}_{\mathcal{Y}_t}\Big[\|\hat{x}^o_{t+N_c}\| - \|\hat{x}^o_t\|\Big] \leqslant \mathbb{E}_{\mathcal{Y}_t}\Big[\|A_o^{N_c}\hat{x}^o_t + \mathfrak{R}_{N_c}(A_o,B_o)u_{t,t+N_c-1}\| - \|\hat{x}^o_t\|\Big] + \zeta \leqslant -\frac{\varepsilon}{2}$ whenever $\|\hat{x}^o_t\| \geqslant \zeta + \varepsilon$ for all $t = T, T+N_c, T+2N_c, \ldots$, where the last inequality follows from the satisfaction of the stability constraint (22); condition (39) of Proposition 11 is satisfied. By orthogonality of A_o, we have that $\|\hat{x}^o_t\| = \|A_o^{N_c}\hat{x}^o_t\|$. It follows by the triangle inequality that, for all

$$\begin{aligned}
t = T, T+N_c, T+2N_c, \ldots\ & \mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}} \\
&\times\Big[\big(\|\hat{x}^o_{t+N_c}\| - \|\hat{x}^o_t\|\big)^4\Big] \\
&= \mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}}\Big[\big(\|\hat{x}^o_{t+N_c}\| - \|A_o^{N_c}\hat{x}^o_t\|\big)^4\Big] \\
&\leqslant \mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}} \\
&\times\Big[\Big(\big\|\mathfrak{R}_{N_c}(A_o,B_o)u_{t,t+N_c-1}\big\| + \big\|(\Xi_t)_{n-n_o+1:n}\big\|\Big)^4\Big].
\end{aligned}$$

Using the inequality $(a+b)^2 \leqslant 2a^2 + 2b^2$, it follows that

$$\begin{aligned}
&\mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}}\Big[\big(\|\hat{x}^o_{t+N_c}\| - \|\hat{x}^o_t\|\big)^4\Big] \\
&\leqslant 8\mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}} \\
&\times\Big[\big\|\mathfrak{R}_{N_c}(A_o,B_o)u_{t,t+N_c-1}\big\|^4 + \big\|(\Xi_t)_{n-n_o+1:n}\big\|^4\Big].
\end{aligned}$$

By design, $\|u_i\|_\infty \leqslant U_{\max}$. In addition, Ξ_t is independent of $\big(\|\hat{x}^o_i\|\big)_{i=T,T+N_c,\ldots,t}$ and is Gaussian; it has its fourth moment bounded. Therefore, there exists a constant $M > 0$ such that $\mathbb{E}_{\{(\|\hat{x}^o_i\|)_{i=T,T+N_c,\ldots,t}\}}[\|\mathfrak{R}_{N_c}(A_o,B_o)u_{t,t+N_c-1}\|^4 + \|(\Xi_t)_{n-n_o+1:n}\|^4] \leqslant M$, for all $t \geqslant T$. The two conditions of Proposition 11 are verified for the sequence $(\|\hat{x}^o_t\|)_{t=T,T+N_c,T+2N_c,\ldots}$. Thus, by Proposition 11, there exists a constant $\gamma'_o > 0$, depending on the given problem parameters, such that $\mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}^o_t\|^2\big] \leqslant \gamma'_o$, for all $t = T, T+N_c, T+2N_c, \ldots$. Finally, using a standard argument (as in Ramponi et al. (2010)) we can show the existence of another constant $\gamma_o \geqslant \gamma'_o$ such that the condition (41) holds. □

Proof (*Proof of* Theorem 5)**.** Claim (i) of Theorem 5 was proved in Lemma 9. It remains to show claim (ii). We start by asserting the following inequality $\mathbb{E}_{\mathcal{Y}_T}\big[\|x_t\|^2\big] \leqslant 2\mathbb{E}_{\mathcal{Y}_T}\big[\|x_t - \hat{x}_t\|^2\big] + 2\mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}_t\|^2\big]$, for all $t \geqslant T = N_c\lceil T'/\kappa\rceil$. We know from Lemma 8 that $\mathbb{E}_{\mathcal{Y}_T}\big[\|x_t - \hat{x}_t\|^2\big] \leqslant \rho$ for all $t \geqslant T \geqslant T'$. We have that $\|\hat{x}_t\|^2 = \|\hat{x}^s_t\|^2 + \|\hat{x}^o_t\|^2$, where $\hat{x}^s$ and $\hat{x}^o$ are states corresponding to the Schur and orthogonal parts of the system, respectively. It then follows that $\mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}_t\|^2\big] = \mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}^s_t\|^2\big] + \mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}^o_t\|^2\big]$. Using Lemmas 10 and 12 we conclude that $\mathbb{E}_{\mathcal{Y}_T}\big[\|\hat{x}_t\|^2\big] \leqslant \gamma_s + \gamma_o$, for all $t \geqslant T$, and therefore, $\mathbb{E}_{\mathcal{Y}_T}\big[\|x_t\|^2\big] \leqslant 2\rho + 2(\gamma_s + \gamma_o)$, for all $t \geqslant T$. Since the sequence $(x_t)_{t\in\mathbb{N}}$ in (1a) is generated through the addition of independent mean-square bounded random variables and a bounded control input, and since $T < \infty$, it follows that there exists a constant $\gamma > 0$ such that $\mathbb{E}_{\mathcal{Y}_0}\left[\|x_t\|^2\right] \leqslant \gamma$, for all $t \in \mathbb{N}$, establishing the claim (ii) of Theorem 5. $\square$

Proof (*Proof of* Corollary 6)**.** The proof of Corollary 6 follows exactly the same reasoning as in the proof of Theorem 5, except for the constraints in (34) and (35). Rewriting the constraints (30) and (31) as (34) and (35), respectively, can be done similarly to the way we rewrote the cost (6) in Theorem 5. It remains to show the upper bounds α^* and β^* in (33). For the constraint (30) we have the upper-bound $\mathbb{E}_{\mathcal{Y}_t}\left[\|X_t\|_{\mathcal{S}}^2 + \mathcal{L}^\top X_t\right] = \mathbb{E}_{\mathcal{Y}_t}\left[\|\mathcal{A}x_t + \mathcal{B}U_t + \mathcal{D}W_t\|_{\mathcal{S}}^2 + \mathcal{L}^\top(\mathcal{A}x_t + \mathcal{B}U_t + \mathcal{D}W_t)\right] \leqslant 3\mathbb{E}_{\mathcal{Y}_t}\left[\|\mathcal{A}x_t\|_{\mathcal{S}}^2 + \|\mathcal{B}U_t\|_{\mathcal{S}}^2 + \|\mathcal{D}W_t\|_{\mathcal{S}}^2\right] + \mathcal{L}^\top\mathcal{A}\hat{x}_t + |\mathcal{L}^\top\mathcal{B}U_t| \leqslant \mathcal{L}^\top\mathcal{A}\hat{x}_t + 3Nm\sigma_{\max}(\mathcal{B}^\top\mathcal{S}\mathcal{B})U_{\max}^2 + \left\|\mathcal{L}^\top\mathcal{B}\right\|_1 U_{\max} + 3\,\mathrm{tr}(\mathcal{A}^\top\mathcal{S}\mathcal{A}\mathbb{E}_{\mathcal{Y}_t}\left[x_t x_t^\top\right] + \mathcal{D}^\top\mathcal{S}\mathcal{D}\Sigma_w)$, where the first inequality follows from the fact that $(a + b + c)^2 \leqslant 3(a^2 + b^2 + c^2)$, for any $a, b, c > 0$, and the noise being zero-mean; the second inequality follows from applying norm bounds between the 2 and ∞-norms and Hölder's inequality. As for the constraint (31), we have the bound $\mathbb{E}_{\mathcal{Y}_t}\left[U_t^\top\tilde{\mathcal{S}}U_t\right] \leqslant \sigma_{\max}(\tilde{\mathcal{S}})\,\|U_t\|^2 \leqslant Nm\sigma_{\max}(\tilde{\mathcal{S}})\,\|U_t\|_\infty^2 \leqslant Nm\sigma_{\max}(\tilde{\mathcal{S}})U_{\max}^2$. $\square$

References

Batina, I. Model predictive control for stochastic systems by randomized algorithms. *Ph.D. thesis*, Technische Universiteit Eindhoven, 2004.

Ben-Tal, A., Boyd, S., & Nemirovski, A. (2006). Extending scope of robust optimization: Comprehensive robust counterparts of uncertain problems. *Journal of Mathematical Programming*, *107*, 63–89.

Ben-Tal, A., Goryashko, A., Guslitzer, E., & Nemirovski, A. (2004). Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, *99*(2), 351–376.

Ben-Tal, A., & Nemirovski, A. (2001). *Lectures on modern convex optimization: analysis, algorithms, and engineering Practice*. SIAM.

Bernardini, D., & Bemporad, A. (2009) Scenario-based model predictive control of stochastic constrained linear systems. In *IEEE conference on decision and control* (pp. 6333–6338). Shanghai, China.

Bernstein, D. S. (2009). *Matrix mathematics* (2nd ed.) Princeton University Press.

Bertsekas, D. P. (2000). *Dynamic programming and optimal control, Vol. 1* (2nd ed.) Athena Scientific.

Bertsekas, D. P. (2007). *Dynamic programming and optimal control, Vol. 2* (3rd ed.) Athena Scientific.

Bertsimas, D., & Brown, D. B. (2007). Constrained stochastic LQC: a tractable approach. *IEEE Transactions on Automatic Control*, *52*(10), 1826–1841.

Bertsimas, D., Iancu, D. A., & Parrilo, P. A. (2010). Optimality of affine policies in multi-stage robust optimization. *Mathematics of Operations Research*, *35*(2), 363–394.

Blackmore, L. A probabilistic particle control approach to optimal, robust predictive control. In *Proceedings of the AIAA guidance, navigation and control conference*, Aug. 2006.

Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. Cambridge: Cambridge University Press, Sixth printing with corrections, 2008.

Cannon, M., Cheng, Q., Kouvaritakis, B., & Raković, S.V. Stochastic tube MPC with state estimation. In *Proc. of the 19th international symposium on mathematical theory of networks and systems–MTNS 2010* (pp. 11–16), July 2010.

Cannon, M., Kouvaritakis, B., & Wu, X. (2009a). Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, *45*(1), 167–172.

Cannon, M., Kouvaritakis, B., & Wu, X. (2009b). Probabilistic constrained MPC for systems with multiplicative and additive stochastic uncertainty. *IEEE Transactions on Automatic Control*, *54*(7), 1626–1632.

Chatterjee, D., Hokayem, P., & Lygeros, J. Stochastic receding horizon control with bounded control inputs: a vector space approach. IEEE Transactions on Automatic Control, 2009. http://arxiv.org/abs/0903.5444 (in press).

Eugenio, E., Agarwal, M., Chatterjee, D., & Lygeros, J. (2011). Convexity and convex approximations of discrete-time stochastic control problems with constraints. *Automatica*, *47*(9), 2082–2087.

Couchman, P. D., Cannon, M., & Kouvaritakis, B. (2006). Stochastic MPC with inequality stability constraints. *Automatica*, *42*(12), 2169–2174.

Goulart, P. J., Kerrigan, E. C., & Maciejowski, J. M. (2006). Optimization over state feedback policies for robust control with constraints. *Automatica*, *42*(4), 523–533.

Hokayem, P., Chatterjee, D., & Lygeros, J. (2009). On stochastic model predictive control with bounded control inputs. In *Proceedings of the combined 48th IEEE conference on decision & control and 28th Chinese control conference* (pp. 6359–6364). http://arxiv.org/abs/0902.3944.

Hokayem, P., Chatterjee, D., Ramponi, F., Chaloulos, G., & Lygeros, J. Stable stochastic receding horizon control of linear systems with bounded control inputs. In *Int. symposium on mathematical theory of networks and systems (MTNS)* (pp. 31–36, Budapest). Hungary, July 2010.

Hokayem, P., Cinquemani, E., Chatterjee, D., & Lygeros, J. Stochastic MPC with imperfect state information and bounded controls. In *Proceedings of the UKACC int. conference on control*. Coventry, UK, Sept. 2010.

Horn, R. A., & Johnson, C. R. (1990). *Matrix analysis*. Cambridge: Cambridge University Press.

Jazwinski, A. H. (1970). *Stochastic processes and filtering theory*. Academic Press.

Kamen, E. W., & Su, J. K. (1999). *Introduction to optimal estimation*. London, UK: Springer.

Klein Haneveld, W. K. (1983). On integrated chance constraints. In *Lecture notes in control and inform. sci.*: *Vol. 76*. *Stochastic programming (Gargnano)* (pp. 194–209). Berlin: Springer.

Klein Haneveld, W. K., & van der Vlerk, M. H. (2006). Integrated chance constraints: reduced forms and an algorithm. *Computational Management Science*, *3*(4), 245–269.

Kumar, P. R., & Varaiya, P. (1986). *Stochastic systems: estimation, identification, and adaptive control*. Prentice Hall.

Löfberg, J. Minimax approaches to robust model predictive control. *Ph.D thesis*, Linköpings Universitet, 2003.

Löfberg, J. (2004). YALMIP: a toolbox for modeling and optimization in MATLAB, In *Proceedings of the CACSD conference*, Taipei, Taiwan.

Li, P., Wendt, M., & Wozny, G. (2002). A probabilistically constrained model predictive controller. *Automatica*, *38*(7), 1171–1176.

Maciejowski, J. M. (2001). *Predictive control with constraints*. Prentice Hall.

Maciejowski, M., Lecchini, A., & Lygeros, J. (2007). NMPC for complex stochastic systems using Markov Chain Monte Carlo. In *Lecture notes in control and information sciences*: *Vol. 358*. *Assessment and future directions of nonlinear model predictive control* (pp. 269–281). Stuttgart, Germany: Springer.

Magni, L., Pala, D., & Scattolini, R. Stochastic model predictive control of constrained linear systems with additive uncertainty. In *European control conference* (pp. 2235–2240). Budapest, Hungary, 2009.

Mayne, D. Q., Rawlings, J. B., Rao, C. V., & Scokaert, P. O. M. (2000). Constrained model predictive control: stability and optimality. *Automatica*, *36*(6), 789–814.

Oldewurtel, F., Jones, C.N., & Morari, M. A tractable approximation of chance constrained stochastic MPC based on affine disturbance feedback. In *Proceedings of the IEEE conference on decision and control*, Cancun, Mexico, Dec. 2008.

Pemantle, R., & Rosenthal, J. S. (1999). Moment conditions for a sequence with negative drift to be uniformly bounded in L^r. *Stochastic Processes and their Applications*, *82*(1), 143–155.

Primbs, J. A., & Sung, C. H. (2009). Stochastic receding horizon control of constrained linear systems with state and control multiplicative noise. *IEEE Transactions on Automatic Control*, *54*(2), 221–230.

Ramponi, F., Chatterjee, D., Milias-Argeitis, A., Hokayem, P., & Lygeros, J. (2010). Attaining mean square boundedness of a marginally stable stochastic linear system with a bounded control input. *IEEE Transactions on Automatic Control*, *55*(10), 2414–2418.

Schwarm, A., & Nikolaou, M. (1999). Chance-constrained model predictive control. *AIChE Journal*, *45*(8), 1743–1752.

Skaf, J., & Boyd, S. (2010). Design of affine controllers via convex optimization. *IEEE Transactions on Automatic Control*, *55*(11), 2476–2487.

Toh, K. C., Todd, M. J., & Tatuncu, R. H. (1999). SDPT3—a Matlab software package for semidefinite programming. *Optimization Methods and Software*, (11), 545–581. http://www.math.nus.edu.sg/mattohkc/sdpt3.html.

van Hessem, D.H., & Bosgra, O.H. A full solution to the constrained stochastic closed-loop MPC problem via state and innovations feedback and its receding horizon implementation, In *Proceedings of the 42nd IEEE conference on decision and control*, vol. 1, (pp. 929–934). 2003.

Wang, Y., & Boyd, S. (2009). Performance bounds for linear stochastic control. *Systems and Control Letters*, *58*(3), 178–182.

Yan, J., & Bitmead, R. R. (2005). Incorporating state estimation into model predictive control and its application to network traffic control. *Automatica*, 595–604.

Yang, Y. D., Sontag, E. D., & Sussmann, H. J. (1997). Global stabilization of linear discrete-time systems with bounded feedback. *Systems and Control Letters*, *30*(5), 273–281.

Peter Hokayem received his B.E. degree in Computer and Communications Engineering from the American University of Beirut in 2001, his M.S. degree in Electrical Engineering from the University of New Mexico in 2003, and his Ph.D. degree in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2007. Following a stint at the University of Wuerzburg, he joined the Automatic Control Laboratory at ETH Zurich in 2008, where he is currently a senior postdoctoral researcher. His current research interests include stochastic model predictive control and distributed optimization and control. He received the Automatica Paper Prize Award in 2008.

Eugenio Cinquemani received his Ph.D. degree in 2005 from the Department of Information Engineering of the University of Padova, Italy. He has then been working as a Post-Doc at the Automatic Control Laboratory of the ETH Zurich, Switzerland, for three years. Since November 2009 he works as a Research Scientist at the INRIA Grenoble — Rhône-Alpes, France. His research interests include identification, estimation and control methods for stochastic and hybrid systems, with applications to biological networks.

Debasish Chatterjee received his Ph.D. in Electrical Engineering from University of Illinois at Urbana-Champaign, USA, in 2007, following which he did his postdoctoral research at the Automatic Control Laboratory, ETH Zurich, Switzerland. He is now with the Systems & Control Engineering, IIT-Bombay, India. His research interests center around constrained and optimization-based control, switched and hybrid systems, and applications of random processes in systems and control.

Federico Ramponi was born in Verona, Italy, on August 3, 1978. He received his M.Sc. in computer engineering and his Ph.D. in control engineering from the University of Padova, Italy, in 2004 and 2009 respectively. From 2009 to 2011 he has been a post-doctoral Fellow within the Automatic Control Laboratory, ETH Zurich, Zurich, Switzerland. He has recently become an assistant professor at the Department of Information Engineering, University of Brescia, Italy. His research interests include stochastic model predictive control and stochastic model checking.

John Lygeros completed a B.Eng. degree in Electrical Engineering in 1990 and an M.Sc. degree in Systems Control in 1991, both at Imperial College, London, UK. In 1996 he obtained a Ph.D. in Electrical Engineering and Computer Sciences from the University of California, Berkeley. During the period 1996–2000 he held postdoctoral research appointments at MIT and UC Berkeley, as well as a part-time research engineer position at SRI International, Menlo Park, California. Between 2000 and 2003 he was a University Lecturer at the Department of Engineering, University of Cambridge, UK, and a Fellow of Churchill College. Between 2003 and 2006 he was an Assistant Professor at the Department of Electrical and Computer Engineering, University of Patras, Greece. In July 2006 he joined the Automatic Control Laboratory at ETH Zurich, where he is currently serving as a Full Professor and the Head of the laboratory. His research interests include modeling, analysis, and control of hierarchical, hybrid, and stochastic systems, with applications to biochemical networks, air traffic management, power grids and camera networks. John Lygeros is a Fellow of the IEEE, and a member of the IET and the Technical Chamber of Greece.