Stochastic Tube MPC with State Estimation

Mark Cannon, **Qifeng Cheng**, **Basil Kouvaritakis**, **Saša V. Raković**

Department of Engineering Science, University of Oxford, UK

Abstract: An output feedback Model Predictive Control (MPC) strategy for linear systems with additive stochastic disturbances and probabilistic constraints is proposed. Given the probability distributions of the disturbance input, the measurement noise and the initial state estimation error, the distributions of future realizations of the constrained variables are predicted using the dynamics of the plant and a linear state estimator. From these distributions, a set of deterministic constraints are computed for the predictions of a nominal model. The constraints are incorporated in a receding horizon optimization of an expected quadratic cost, which is formulated as a quadratic program. The constraints are constructed so as to provide a guarantee of recursive feasibility, and the closed loop system is stable in a mean-square sense.

Keywords: Output feedback; probabilistic constraints; stochastic systems

1. INTRODUCTION

Model predictive control (MPC) strategies have proved highly successful as a result of their ability to achieve approximately optimal performance in the presence of constraints. The vast majority of contributions in this area have considered hard constraints, and a number of robust MPC techniques have been proposed for handling hard constraints in the case of systems that are subject to uncertainty. However model and measurement uncertainties are often stochastic, and in such cases the robust MPC approach can be conservative because it ignores information on the probabilistic distribution of the uncertainty. In addition, not all constraints are hard, and it may be possible, as well as desirable in order to improve performance, to tolerate violations of some constraints provided that the frequency of violations remains within allowable limits.

Recent work (Cannon et al., 2010; Kouvaritakis et al., 2010) proposed Stochastic MPC (SMPC) algorithms that made explicit use of probabilistic information on additive disturbances in order to minimize (in a receding horizon fashion) the expected value of a predicted cost subject to a combination of hard and soft constraints. A key ingredient of the algorithms is the definition of stochastic tubes that enable a recursive guarantee of feasibility and therefore enable the assertion of closed loop stability and constraint satisfaction. Probabilistic constraints were invoked via linear constraints on nominal predictions using the concept of tubes (Langson et al., 2004; Mayne et al., 2005, 2006), and thus the online algorithms are computationally efficient. Moreover Kouvaritakis et al. (2010) proposed conditions that are necessary as well as sufficient for the probabilistic constraints and their recursive feasibility guarantee, implying minimal conservativism in the handling of probabilistic constraints. The approach was based on state feedback, which carries with it the assumption that the states are measurable. In practice this is typically not the case, and it is often necessary therefore to estimate the state via an observer.

The mechanism for incorporating state estimation in robust MPC is well understood (see e.g. Lee and Kouvaritakis, 2001; Wan and Kothare, 2002; Løvaas et al., 2008; Mayne et al., 2009) and uses lifting in order to get a description of the combined dynamics of the system and observer. However it is now necessary to integrate into the approach the probabilistic information that is usually available on the measurement noise and the distribution of the unknown initial plant state variables. This paper considers the propagation of this information, together with information on the distribution of an additive disturbance, through the prediction dynamics. Whereas previous work on this problem (e.g. Yan and Bitmead, 2005) did not consider feasibility and stability, the approach proposed here guarantees recursive feasibility with respect to both hard and probabilistic constraints and ensures stability and convergence of the plant state in a mean-square sense. We provide a numerical example verifying that the rate of constraint violation can be as tight as the specified probability.

2. PROBLEM FORMULATION

Consider a linear system with state $x_k \in \mathbb{R}^n$, measured output $y_k \in \mathbb{R}^{n_y}$ and control input $u_k \in \mathbb{R}^{n_u}$:

$$x_{k+1} = Ax_k + Bu_k + Dw_k \tag{1a}$$
$$y_k = Cx_k + Fv_k \tag{1b}$$

where the disturbance $w_k \in \mathbb{R}^{n_w}$ and measurement noise $v_k \in \mathbb{R}^{n_v}$ sequences are independent, identically distributed, with known, finitely supported distributions such that the elements of w_k and v_k are zero-mean and independent. Consider also the linear observer dynamics:

$$\hat{x}_{k+1} = A\hat{x}_k + Bu_k + L(y_k - \hat{y}_k) \tag{2a}$$
$$\hat{y}_k = C\hat{x}_k \tag{2b}$$

ISBN 978-963-311-370-7

where $\hat{x}_k$ is the state estimate and $L \in \mathbb{R}^{n \times n_y}$ is chosen such that $A - LC$ is strictly stable. Subtracting (2a) from (1a), the error dynamics are given by

$$\epsilon_{k+1} = (A - LC)\epsilon_k + Dw_k + Fv_k \tag{3}$$

where $\epsilon_k := x_k - \hat{x}_k$ is the state estimation error. The initial value ϵ_0 is assumed to be a random variable with a known and compactly supported distribution, so that $\epsilon_0 \in \Pi_0$ for some compact set Π_0.

We denote the state estimate and estimation error sequences predicted at time k as $\{\hat{x}_{k+i|k}, \epsilon_{k+i|k},\ i = 0, 1, \ldots\}$. Using the closed loop paradigm (Kouvaritakis et al., 2000), the augmented system state $\xi_{k+i} = (\hat{x}_{k+i}, \epsilon_{k+i}) \in \mathbb{R}^{2n}$ evolves according to

$$\xi_{k+i+1|k} = \Psi \xi_{k+i|k} + \tilde{B} c_{k+i|k} + \tilde{D}\delta_{k+i|k} \tag{4a}$$

$$u_{k+i|k} = K\hat{x}_{k+i|k} + c_{k+i|k} \tag{4b}$$

where $\delta_{k+i|k} = (w_{k+i|k}, \nu_{k+i|k}) \in \mathbb{R}^{n_w+n_v}$ is the stochastic disturbance, Ψ is assumed to be strictly stable, and

$$\Psi = \begin{bmatrix} A+BK & LC \\ 0 & A-LC \end{bmatrix}, \quad \tilde{B} = \begin{bmatrix} B \\ 0 \end{bmatrix}, \quad \tilde{D} = \begin{bmatrix} 0 & LF \\ D & -LF \end{bmatrix}.$$

In this formulation $c_{k+i|k}$, $i = 0, \ldots, N-1$ are decision variables, with $c_{k+i|k} = 0$ for $i \geq N$, for some finite prediction horizon N. Given the assumptions on w_k and v_k, the distribution of δ_k is compactly supported, with $\delta_k \in \Delta$ for all k, where Δ is a compact set.

The system is subject to probabilistic constraints on linear functions of the state and input of the form

$$\Pr\{\eta_x^T x_k + \eta_u^T u_k \leq h\} \geq p, \quad k = 0, 1, \ldots \tag{5}$$

for fixed scalar h, vectors η_x, η_u, and probability $p \in (0, 1]$. Applied to the predictions of (4), these are equivalent to

$$\Pr\big\{g^T \xi_{k+i|k} + f^T c_{k+i|k} \leq h\big\} \geq p, \quad i = 0, 1, \ldots \tag{6}$$

where $g = (\eta_x + \eta_u K, \eta_x)$ and $f = \eta_u$. Note that (6) applies for $i = 1, 2, \ldots$ if $\eta_u = 0$. The case of hard constraints is treated simply by setting $p = 1$ in (5) and (6).

We split the prediction $\xi_{k+i|k}$ into nominal ($z_{k+i|k}$) and uncertain ($e_{k+i|k}$) elements:

$$\xi_{k+i|k} = z_{k+i|k} + e_{k+i|k} \tag{7a}$$

$$z_{k+i+1|k} = \Psi z_{k+i|k} + \tilde{B} c_{k+i|k} \tag{7b}$$

$$e_{k+i+1|k} = \Psi e_{k+i|k} + \tilde{D}\delta_{k+i|k} \tag{7c}$$

where $z_{k|k} = (\hat{x}_k, 0)$, $e_{k|k} = (0, \epsilon_k)$, $e_{k|k} \sim \mathfrak{D}_k$, and the distribution $\mathfrak{D}_k$ can be calculated according to (3) given the distribution of the initial estimation error ϵ_0. This decomposition allows the propagation of the disturbance, $e_{k+i|k}$, to be considered separately from the nominal state prediction, $z_{k+i|k}$, and thus simplifies the handling of constraints.

The problem is then to devise a receding horizon MPC strategy that minimizes the cost

$$J_k = \sum_{i=0}^{\infty} \mathbb{E}\left(x_{k+i|k}^T Q x_{k+i|k} + u_{k+i|k}^T R u_{k+i|k}\right) \tag{8}$$

(where $\mathbb{E}(\cdot)$ denotes expectation) subject to satisfaction of the constraint (5) for all $k \geq 0$, while ensuring that the closed loop system stable and that x_k converges to a neighbourhood of the origin.

3. PREDICTIONS AND RECURSIVELY FEASIBLE PROBABILISTIC TUBES

To handle the constraints of (5), we first consider the conditions that guarantee that (6) is satisfied by predicted state and input sequences. In the sequel, the vector of decision variables at each time k is denoted as $\mathbf{c}_k = (c_{k|k}, c_{k+1|k}, \ldots, c_{k+N-1|k}) \in \mathbb{R}^{Nn_u}$. Necessary and sufficient conditions are given as follows.

Theorem 1. At time k, the constraints (6) are satisfied by predictions of (4) if and only if $\mathbf{c}_k$ satisfies

$$(g^T H_i + f^T E_i)\mathbf{c}_k + g^T \Psi^i z_k \leq h - \hat{\gamma}_{i|k}, \quad i = 0, 1, \ldots \tag{9}$$

where $H_i = [\Psi^{i-1}\tilde{B}\ \cdots\ \tilde{B}\ 0\ \cdots\ 0]$, $E_i \mathbf{c}_k = c_{k+i|k}$, and $\hat{\gamma}_{i|k}$ is defined for each k and $i = 0, 1, \ldots$ as the minimum value such that

$$\Pr\big\{g^T e_{k|k} \leq \hat{\gamma}_{0|k}\big\} = p \tag{10a}$$

$$\Pr\big\{g^T(\Psi^i e_{k|k} + \Psi^{i-1}\tilde{D}\delta_{k|k} + \cdots + \tilde{D}\delta_{k+i-1|k}) \leq \hat{\gamma}_{i|k}\big\} = p, \quad \text{for } i = 1, 2, \ldots \tag{10b}$$

Proof: The predictions of (7) give

$$z_{k+i|k} = \Psi^i z_{k|k} + H_i \mathbf{c}_k \tag{11a}$$

$$e_{k+i|k} = \Psi^i e_{k|k} + \Psi^{i-1}\tilde{D}\delta_{k|k} + \cdots + \tilde{D}\delta_{k+i-1|k} \tag{11b}$$

and since $\hat{\gamma}_{i|k}$ is the minimum value that satisfies (10), it follows that (9) is equivalent to the constraint of (6). ■

Many problems of practical interest have more than one constrained variable. For the case of r constrained scalar variables, we assume that each constraint can be expressed as a condition of the form of (5), and hence (6), i.e.

$$\Pr\big\{g_j^T \xi + f_j^T c \leq h_j\big\} \geq p_j,$$

for $j = 1, \ldots, r$. The treatment of Theorem 1 can then be applied to each of the r individual constraints to derive a set of constraints of the form (9).

Condition (9) is linear in $\mathbf{c}_k$, and can therefore be conveniently incorporated into an online optimization of $\mathbf{c}_k$. However the computation of $\hat{\gamma}_{i|k}$ necessitates calculating the distribution of $g^T(\Psi^i e_{k|k} + \Psi^{i-1}\tilde{D}\delta_k + \cdots + \tilde{D}\delta_{k+i-1})$, which involves a multivariate convolution integral, and hence can be computationally intensive, particularly for high-dimensional systems and long prediction horizons. However, by discretizing the distributions of δ_k and ϵ_0 and performing discrete scalar convolutions, the values of $\hat{\gamma}_{i|k}$ can be approximated (to a specified degree of accuracy) at reasonable computational cost. For example, computing the discrete density function of a sum of two scalar variables requires $O(m^2)$ multiplications if each variable can take m discrete values. Most importantly, the calculation of $\hat{\gamma}_{i|k}$ does not require knowledge of x_k and can therefore be performed offline.

Although the conditions of Theorem 1 ensure that (6) is satisfied over the entire prediction horizon at time k, the existence of $\mathbf{c}_k$ satisfying these conditions does not ensure the existence of a set of decision variables $\mathbf{c}_{k+1}$ that generate predictions at time $k+1$ satisfying (6). Hence (9) does not guarantee the future feasibility of an online optimization problem incorporating (9) as a constraint. This is because $\hat{\gamma}_{i|k}$ in (10) is a probabilistic bound on the stochastic components of the predicted value of

$g^T\xi_{k+i|k}$ at time k, whereas the same probabilistic bound on $g^T\xi_{k+i|k+1}$ could, on the basis of information available at time k, take its maximum value over all possible realizations of the estimation error e_k and disturbance δ_k at time k.

Define $T\mathbf{c}_k = (c_{k+1|k}, \dots, c_{k+N-1|k}, 0) \in \mathbb{R}^{Nn_u}$, then the preceding argument implies that $\mathbf{c}_{k+1} = T\mathbf{c}_k$ satisfies (9) at time $k+1$ if and only if $\hat{\gamma}_{i|k}$ is replaced in (9) by the bounds that are obtained by assuming the worst-case values for $e_{k|k}$ and $\delta_{k|k}$ in (10). The following theorem generalizes this approach by deriving the conditions for feasibility of $\mathbf{c}_{k+i} = T^i\mathbf{c}_k$ in (9) at time $k+i$ for $i \geq 1$. These conditions therefore ensure that the probabilistic constraints (6) are satisfied at time k and are also recursively feasible in the sense that they remain feasible at all times $k+1, k+2, \dots$ if feasible at time k.

In the sequel, γ_i is defined as the minimum value such that

$$\Pr\big\{g^T(\Psi^{i-1}\tilde{D}\delta_k + \dots + \tilde{D}\delta_{k+i-1}) \leq \gamma_i\big\} = p. \tag{12}$$

We also define $\alpha_{i|k}$ and d_i for $i = 0, 1, \dots$ by

$$d_i = \max_{\delta\in\Delta}\; g^T\Psi^{i-1}\tilde{D}\delta \tag{13a}$$

$$\alpha_{i|k} = \max_{e_{k|k}\sim\mathfrak{D}_k} g^T\Psi^i e_{k|k} \tag{13b}$$

where $\max_{e\sim\mathfrak{D}}\{\cdot\}$ denotes the maximum over all realizations of e in the distribution defined by $\mathfrak{D}$.

Theorem 2. At time k, the constraints (6) are satisfied by the predictions of (4) if $\mathbf{c}_k$ satisfies

$$(g^TH_i+f^TE_i)\mathbf{c}_k+g^T\Psi^i z_{k|k} \leq h-\beta_{i|k}, \quad i = 0, 1, \dots \tag{14}$$

where $\beta_{i|k}$ is the maximum element of the $(i+1)$th column of the matrix $\mathfrak{B}_k$ defined by

$$\begin{bmatrix} \hat{\gamma}_{0|k} & \hat{\gamma}_{1|k} & \hat{\gamma}_{2|k} & \hat{\gamma}_{3|k} & \cdots \\ 0 & \alpha_{1|k}+d_1 & \alpha_{2|k}+d_2+\gamma_1 & \alpha_{3|k}+d_3+\gamma_2 & \cdots \\ 0 & 0 & \alpha_{2|k}+d_2+d_1 & \alpha_{3|k}+d_3+d_2+\gamma_1 & \cdots \\ 0 & 0 & 0 & \alpha_{3|k}+d_3+d_2+d_1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \end{bmatrix} \tag{15}$$

Furthermore, if (14) is feasible at time k, then (14) will remain feasible at all times $k+1, k+2, \dots$.

Proof: If $\beta_{i|k}$ is defined by the $(i+1)$th element of the first row of $\mathfrak{B}_k$, then (14) is equivalent to (9). The $(i+1)$th element of the second row of $\mathfrak{B}_k$ is obtained by replacing $e_{k|k}$ and $\delta_{k|k}$ in (10b) by their worst-case values, so that (14) also ensures that $T\mathbf{c}_k$ is feasible for (9) at time $k+1$. Similarly the jth row corresponds to the worst-case values of $e_{k|k}$ and $\delta_{k|k}, \dots, \delta_{k+j-2|k}$ in (10b), thus ensuring that $T^j\mathbf{c}_k$ is feasible for (9) at time $k+j$. The proof is completed by noting that the conditions of (14) are themselves recursively feasible due to their definition in terms of the worst-case future uncertainty. ∎

Remark 3. If $\eta_u = 0$ in (5), so that the system constraints apply only to the plant state, then the constraint (6) and condition (14) should be invoked for $i = 1, 2, \dots$, so that the first prediction time-step ($i = 0$) is excluded. In this case $\mathfrak{B}_k$ is given by the matrix of (15) with the elements on the main diagonal set to zero.

Lemma 4. For all $i \geq 1$ we have $\beta_{i|k} = \alpha_{i|k} + \sum_{j=1}^{i} d_j$.

Proof: From the definitions of d_i and $\alpha_{i|k}$ in (13) we have $g^T\Psi^{i-1}\tilde{D}\delta_k \leq d_i$ and $g^T\Psi^i e_{k|k} \leq \alpha_{i|k}$ for all realizations of $e_{k|k}$ and δ_k. Hence $\hat{\gamma}_{1|k} \leq \alpha_{1|k}+d_1$, and from the definition of γ_{i-1} it also follows that

$$\Pr\left\{\begin{aligned} g^T(\Psi^i e_{k|k} + \Psi^{i-1}\tilde{D}\delta_k + \dots + \tilde{D}\delta_{k+i-1}) \\ \leq \alpha_{i|k} + d_i + \gamma_{i-1} \end{aligned}\right\} \geq p$$

so that $\hat{\gamma}_{i|k} \leq \alpha_{i|k} + d_i + \gamma_{i-1}$ for $i = 2, 3, \dots$. Similar arguments give $\gamma_1 \leq d_1$ and $\gamma_i \leq d_i + \gamma_{i-1}$ for $i = 2, 3, \dots$, and it follows that the maximum element in each column of $\mathfrak{B}_k$ lies on the diagonal. ∎

Since $e_{k|k}$ and the sequence $\{\delta_k, \delta_{k+1}, \dots\}$ are (by assumption) bounded, and since Ψ is strictly stable, it is clear that the predicted sequence $\{e_{k|k}, e_{k+1|k}, \dots\}$ generated by (7c) is bounded, and hence the sequence $\{\beta_{i|k}, \beta_{i+1|k} \dots\}$ is bounded for any k. We next show how to compute bounds on $\beta_{i|k}$ that hold for all i, k exceeding given values. These are used in section 4 to derive a recursively feasible SMPC algorithm that involves only a finite number of constraints. As we show below, these bounds also enable the constraints in the online SMPC optimization at all times $k = 0, 1, \dots$ to be computed offline, at $k = 0$.

Define the symmetric positive definite matrix $S \in \mathbb{R}^{2n\times 2n}$ and scalar $\rho \in (0, 1)$ as the solution of the semidefinite program:

$$(\rho, S) = \arg \min_{\substack{\rho>0 \\ S=S^T\succ 0}} \; \rho$$

$$\text{subject to} \quad \Psi S\Psi^T \preceq \rho^2 S \tag{16a}$$

$$\|\tilde{D}\delta\|_{S^{-1}} \leq 1 \quad \text{for all } \delta \in \Delta \tag{16b}$$

(where $\|v\|_S = \sqrt{v^TSv}$). Define $\mathcal{E}$ as the ellipsoidal set $\mathcal{E} = \{e : e^TS^{-1}e \leq 1\}$, and let $\mathcal{E}_\epsilon$ denote its projection onto the subspace $\{e \in \mathbb{R}^{2n} : [I_n \;\; 0]e = 0\}$, so that

$$\mathcal{E}_\epsilon = \{\epsilon : \epsilon^TS_{22}^{-1}\epsilon \leq 1\},$$

where S_{22} is the relevant block of $S = \begin{bmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{bmatrix}$.

Also define positive scalars τ and σ by

$$\tau = \max_{e_0\sim\mathfrak{D}_0} \|e_0\|_{S^{-1}} = \max_{\epsilon_0\in\Pi_0} \left\|\begin{bmatrix} 0 \\ \epsilon_0 \end{bmatrix}\right\|_{S^{-1}} \tag{17a}$$

$$\sigma = \max_{\epsilon\in\mathcal{E}_\epsilon} \left\|\begin{bmatrix} 0 \\ \epsilon \end{bmatrix}\right\|_{S^{-1}} = \lambda_{\max}\Big(S_{22}^{1/2}\,[0 \;\; I]\, S^{-1} \begin{bmatrix} 0 \\ I \end{bmatrix} S_{22}^{1/2}\Big). \tag{17b}$$

(where $\lambda_{\max}(A)$ denotes the maximum eigenvalue of A).

Lemma 5. For any integer $\nu > 1$, $\beta_{i|k}$ satisfies:

$$\beta_{i|k} = \sum_{j=1}^{i} d_j + \sum_{j=1}^{k} b_{ij} + a_{ik}, \qquad i = 1, 2, \dots, \nu-1, \; k < \nu \tag{18a}$$

$$\beta_{i|k} \leq \bar{\beta}_{i|\nu} = \sum_{j=1}^{i} d_j + \sum_{j=1}^{\nu-i} b_{ij} + \bar{\alpha}, \qquad i = 1, 2, \dots, \nu-1, \; k \geq \nu \tag{18b}$$

$$\beta_{i|k} \leq \bar{\beta} = \sum_{j=1}^{\nu} d_j + \bar{\alpha}, \; i \geq \nu \tag{18c}$$

where

$$b_{ij} = \max_{\delta\in\Delta} g^T\Psi^i \begin{bmatrix}0 & 0\\ 0 & I\end{bmatrix} \Psi^{j-1}\tilde{D}\delta$$
$$a_{ik} = \max_{e_0\sim\mathfrak{D}_0} g^T\Psi^i \begin{bmatrix}0 & 0\\ 0 & I\end{bmatrix} \Psi^{k}e_0$$
$$\bar{\alpha} = \rho^\nu\sigma\Big(\frac{1}{1-\rho}+\tau\Big)\|g\|_S.$$

Proof: The definitions of $\alpha_{i|k}$ in (13b) and $e_{k|k}$ in Theorem 1 give

$$\alpha_{i|k} = \sum_{j=1}^{k} b_{ij} + a_{ik},$$

which implies that the expression for $\beta_{i|k}$ in (18a) holds for all $i \geq 1$ and $k \geq 0$. To demonstrate the bounds in (18b,c), note that (16b) implies $\tilde{D}\Delta \subseteq \mathcal{E}$ (i.e. $\tilde{D}\delta \in \mathcal{E}$ for all $\delta \in \Delta$), and (16a) therefore implies $\Psi^j\tilde{D}\Delta \subseteq \rho^j\mathcal{E}$. Since $\max_{e\in\mathcal{E}} g^T e = \|g\|_S$, it follows that

$$d_j \leq \rho^{j-1}\|g\|_S.$$

Furthermore, (17b) implies

$$\begin{bmatrix}0 & 0\\ 0 & I\end{bmatrix}\mathcal{E} \subseteq \sigma\mathcal{E},$$

so by the same argument we have

$$b_{ij} \leq \rho^{i+j-1}\sigma\|g\|_S,$$

and similarly, from (17a),

$$a_{ik} \leq \rho^{i+k}\sigma\tau\|g\|_S.$$

The bounds in (18b) and (18c) can be obtained by substituting these bounds into (18a) and making use of the properties that $\rho < 1$ (since Ψ in (16a) is strictly stable), and $\sigma > 1$ (due to the definition of σ in (17b)). ■

Remark 6. The bounds in (18b,c) can be made arbitrarily tight by using a sufficiently large value for ν.

Remark 7. The values of $\beta_{0|k} = \hat{\gamma}_{0|k}$ for $k \geq \nu$ may (since Ψ is strictly stable) be approximated for sufficiently large ν using the asymptotic distribution $\mathfrak{D}_\infty$ of $\lim_{k\to\infty} e_{k|k}$, i.e.

$$\beta_{0|k} \approx \hat{\gamma}_{0|\infty} = \gamma_\infty \ \forall k \geq \nu.$$

Combining this approximation with the bounds on $\beta_{i|k}$ for $i \geq 1$ in Lemma 5 enables recursively feasible constraints to be computed offline for all $k \geq 0$ on the basis of a finite number of bounds, namely

$$\{\hat{\gamma}_{0|k},\ k=0,\ldots,\nu-1\},\ \gamma_\infty,\ \{d_j,\ j=1,\ldots,\nu\},$$
$$\{b_{ij},\ i,j=1,\ldots\nu-1\},\ \{a_{ij},\ i,j=1,\ldots\nu-1\},\ \bar{\alpha}. \tag{19}$$

This approach implicitly assumes that the transient response of the coupled plant and observer that is due to the initial estimation error becomes negligible after ν time steps. For ν less than the duration of these transients, the degree of conservativeness of bounds of Lemma 5 could be non-negligible.

4. SMPC ALGORITHM

This section describes a receding horizon strategy employing the predictions of (4), and establishes its stability and convergences properties in closed loop operation. The strategy is based on an online optimization which minimizes the expected cost (8) over the decision variables $\mathbf{c}_k = (c_{k|k},\ldots,c_{k+N-1|k})$ subject to the constraints (6), which are invoked via the recursively feasible constraints of (14). The predictions of (4b), with $c_{k+i|k} = 0$ for $i \geq N$, imply a dual mode prediction strategy (Mayne et al., 2000), with mode 1 comprising the first N prediction time steps during which the control inputs are determined by the elements of $\mathbf{c}_k$, and mode 2 the subsequent infinite prediction horizon over which the control inputs are given by the fixed feedback law $u_{k+i|k} = Kx_{k+i|k}$.

Constraints are handled explicitly in mode 1 and implicitly in mode 2 through the constraint that the predicted state at the end of the mode 1 horizon should lie in a terminal set. To derive this terminal set we re-write the constraints (14) in mode 2 at time k as

$$g^T\Psi^j z_{k+N|k} \leq h - \beta_{N+j|k}\,,\quad j=0,1,\ldots \tag{20}$$

The maximal admissible set at time k, denoted $\mathcal{S}_{\infty|k}$, is the subset of $\mathbb{R}^n$ containing all $z_{k+N|k}$ such that (20) holds. The following result uses the bounds (18) to establish conditions under which $S_{\infty|k}$ is well-defined.

Lemma 8. If $h \geq \bar{\beta}$ and

$$h \geq \begin{cases}\max\limits_{i=N,\ldots,\nu-1}\beta_{i|k} & \text{if } k<\nu\\ \max\limits_{i=N,\ldots\nu-1}\bar{\beta}_{i|\nu} & \text{otherwise}\end{cases} \tag{21}$$

then $\mathcal{S}_{\infty|k}$ is non-empty.

Proof: Using (18b,c) to bound $\beta_{N+j|k}$ in the inequalities (20) defining $\mathcal{S}_{\infty|k}$, it is easy to show that the conditions on h given in the lemma ensure that $\mathcal{S}_{\infty|k}$ contains the origin, and so is non-empty. ■

The approach of Gilbert and Tan (1991) provides a method of computing an inner approximation of $\mathcal{S}_{\infty|k}$. Consider first the case of $k<\nu$, and define a terminal set $\mathcal{S}_{\nu|k}$ as follows:

$$\begin{aligned}\mathcal{S}_{\nu|k} = \{z :\ & g^T\Psi^{i-N}z \leq h-\beta_{i|k}, && i=N,\ldots,\nu-1\\ & g^T\Psi^{i-N}z \leq h-\bar{\beta}, && i=\nu,\nu+1,\ldots\}\end{aligned} \tag{22}$$

From (18c) it follows that $\mathcal{S}_{\nu|k} \subseteq \mathcal{S}_{\infty|k}$, and furthermore, it can be shown (Gilbert and Tan, 1991) that there exists N_k^* such that $\mathcal{S}_{\nu|k}$ is determined by the finite set of inequalities:

$$\begin{aligned}\mathcal{S}_{\nu|k} = \{z :\ & g^T\Psi^{i-N}z \leq h-\beta_{i|k}, && i=N,\ldots,\nu-1\\ & g^T\Psi^{i-N}z \leq h-\bar{\beta}, && i=\nu,\ldots,\nu+N_k^*\}.\end{aligned} \tag{23}$$

whenever $\mathcal{S}_{\infty|k}$ is bounded. For $k \geq \nu$, the bounds of (18b) enable a time-invariant terminal set to be defined as $\mathcal{S}_{\nu|k} = \mathcal{S}_\nu$, where

$$\begin{aligned}\mathcal{S}_{\nu} = \{z :\ & g^T\Psi^{i-N}z \leq h-\bar{\beta}_{i|\nu}, && i=N,\ldots,\nu-1\\ & g^T\Psi^{i-N}z \leq h-\bar{\beta}, && i=\nu,\ldots,\nu+N_\nu^*\}\end{aligned} \tag{24}$$

for some finite N_ν^*. Thus a finite collection of terminal sets: $\mathcal{S}_{\nu|k}$, for $k=0,\ldots,\nu-1$ and $\mathcal{S}_\nu$ for all $k \geq \nu$ can be used to invoke the mode 2 constraints at any time k. Moreover, for each $k=0,\ldots,\nu$, the parameter N_k^* can be determined by solving a finite number of linear programming problems (see Gilbert and Tan, 1991, for details).

Note that if $\mathcal{S}_{\infty|k}$ is unbounded, then $\mathcal{S}_{\nu|k}$ and $\mathcal{S}_\nu$ can be defined as bounded inner approximations of $\mathcal{S}_{\infty|k}$ by incorporating suitable additional constraints in (22). The approach outlined above can then be used to express $\mathcal{S}_{\nu|k}$ and $\mathcal{S}_\nu$ in terms of a finite number of inequalities as in (23) and (24).

The cost defined by (8) is the summation of the expected value of each stage cost, which is necessarily infinite over the infinite prediction horizon since the asymptotic value:

$$l_{ss} = \lim_{i\to\infty} \mathbb{E}\big(x_{k+i|k}^T Q x_{k+i|k} + u_{k+i|k}^T R u_{k+i|k}\big)$$

is non-zero for the case of persistent additive uncertainty. Therefore we redefine the cost function in terms of the deviation of expected stage cost from this asymptotic limit:

$$\tilde{J}_k = \sum_{i=0}^{\infty}\Big[\mathbb{E}\big(x_{k+i|k}^T Q x_{k+i|k} + u_{k+i|k}^T R u_{k+i|k}\big) - l_{ss}\Big]. \quad (25)$$

This modification provides a finite cost that can be minimized online. Note however that the optimal value of $\mathbf{c}_k$ is unchanged.

Remark 9. Given the distribution of δ, the cost (25) can be expressed as a quadratic function of the degrees of freedom $\mathbf{c}_k$ in the predictions (see e.g. Cannon et al., 2009):

$$\tilde{J}_k(\mathbf{c}_k) = \mathbf{c}_k^T P_{cc}\mathbf{c}_k + 2\hat{x}_k^T P_{xc}^T \mathbf{c}_k + p_k(\hat{x}_k)$$

where $P_{cc} \in \mathbb{R}^{Nn_u\times Nn_u}$, $P_{xc} \in \mathbb{R}^{n\times Nn_u}$ are constants that can be computed offline, and $p_k(\hat{x}_k)$ is independent of $\mathbf{c}_k$. Furthermore, if K is the optimal feedback gain for the case of no constraints, then $P_{xc} = 0$ by construction.

Algorithm 1. (SMPC).
Offline Use the distributions of e_0 and δ_k to compute the parameters (19) that define the recursively feasible probabilistic constraints (6). Determine the constraint checking horizons N_k^*, $k = 0, \dots, \nu$ in the terminal constraint sets defined in (23) and (24).
Online At each time $k = 0, 1, \dots$:

1. Compute the state estimate $\hat{x}_k$ and set $z_{k|k} = (\hat{x}_k, 0)$.
2. Solve the quadratic program (QP):

$$\begin{aligned}\mathbf{c}_k^* = &\arg\min_{\mathbf{c}_k}\ \tilde{J}_k(\mathbf{c}_k)\\ &\text{subject to } (g^T H_i + f^T E_i)\mathbf{c}_k + g^T\Psi^i z_{k|k} \le h - \beta_i,\\ &\qquad\qquad i = 0, 1, \dots, N-1\\ &\qquad z_{k+N|k} \in \mathcal{S}_k\end{aligned}$$

   where $\beta_i = \beta_{i|k}$ and $\mathcal{S}_k = \mathcal{S}_{\nu|k}$ if $k < \nu$
   $\beta_i = \beta_{i|\nu}$ and $\mathcal{S}_k = \mathcal{S}_\nu$ if $k \ge \nu$.
3. Set $u_k = K\hat{x}_{k+i|k} + c_{k|k}^*$.

Theorem 10. Given feasibility at $k = 0$, SMPC remains feasible for all $k > 0$. The closed loop system under SMPC satisfies the probabilistic constraint (5) and the quadratic stability condition:

$$\lim_{r\to\infty}\frac{1}{r}\sum_{k=0}^{r}\mathbb{E}\big(x_k^T Q x_k + u_k^T R u_k\big) \le l_{ss}. \quad (26)$$

Proof: The recurrence of feasibility is due to the definition of the constraints in the online optimization of SMPC, and this ensures that the probabilistic constraints (5) are satisfied in closed loop operation. In addition, Theorem 2 ensures that $T\mathbf{c}_k^*$ is feasible for the online optimization at $k+1$, so that

$$\tilde{J}_k(\mathbf{c}_k^*) = \mathbb{E}_k\big[\tilde{J}_{k+1}(T\mathbf{c}_k^*)\big] + x_k^T Q x_k + u_k^T R u_k - l_{ss}$$

which therefore implies that the optimal cost value satisfies the stochastic Lyapunov-like condition

$$\tilde{J}_k(\mathbf{c}_k^*) \ge \mathbb{E}_k\big[\tilde{J}_{k+1}(\mathbf{c}_{k+1}^*)\big] + x_k^T Q x_k + u_k^T R u_k - l_{ss} \quad (27)$$

Summing (27) over $k = 0, 1, \dots, r$ gives

$$\sum_{k=0}^{r}\Big[\mathbb{E}\big(x_k^T Q x_k + u_k^T R u_k\big) - l_{ss}\Big] \le \tilde{J}_0(\mathbf{c}_0^*) - \mathbb{E}\big[\tilde{J}_{r+1}(\mathbf{c}_{r+1}^*)\big]$$

which implies (26) since $\tilde{J}(\mathbf{c}_0^*)$ is by assumption finite and (27) ensures that $\mathbb{E}\big[\tilde{J}_r(\mathbf{c}_r^*)\big]$ is finite for all r. ■

Corollary 11. If K is the optimal feedback gain for the case of no constraints, then $x_k \in \mathcal{T}$ in the limit as $k \to \infty$ under SMPC, where $\mathcal{T}$ denotes the the minimal robust invariant set for the state of (1a,b) under $u_k = K\hat{x}_k$.

Proof: The quadratic form of $\tilde{J}(\mathbf{c}_k)$ discussed in Remark 9 and the feasibility of $\mathbf{c}_{k+1} = T\mathbf{c}_k$ together imply that

$$\mathbf{c}_k^T P_{cc}\mathbf{c}_k \ge \mathbf{c}_{k+1}^* P_{cc}\mathbf{c}_{k+1}^* - {c_{k|k}^*}^T(R + B^T P_x B)c_{k|k}^*$$

(where P_x is the steady state solution of the Riccati equation associated with the model (1) and cost (8)). From the sum of both sides of this inequality over all $k \ge 0$ and the fact that $R + B^T P_x B$ is positive definite, it follows that $c_{k|k}^* \to 0$ as $k \to \infty$, which implies that x_k converges to $\mathcal{T}$. ■

5. SIMULATION EXAMPLE

In this example, the model parameters in (1) are given by

$$A = \begin{bmatrix}1.6 & 1.1\\ -0.7 & 1.2\end{bmatrix},\ B = \begin{bmatrix}1\\1\end{bmatrix}\ C = [0.9\ \ 0.2],\ D = I,\ F = 1$$

and the probabilistic constraint (5) defines a state constraint, with

$$\eta_x^T = [1\ \ 0.3],\quad \eta_u = 0,\quad h = 4.5,\quad p = 0.8.$$

The quadratic cost weights are $Q = C^T C$ and $R = 1$.

The distributions of the stochastic model parameters are derived from truncated (and appropriately re-scaled) normal distributions. For example, the measurement noise v_k is derived from a normally distributed random variable, $v_k' \sim \mathcal{N}(0, 1/24^2)$, which is truncated so that $|v_k| \le 0.12$. Therefore the distribution function for v_k (defined by $F_v(V) = \Pr(v_k \le V)$) is obtained from that of v_k' by an affine transformation

$$F_v(V) = aF_{v'}(V) + b \quad \text{for all } V \in [-0.12, 0.12],$$

where a, b are such that $F_v(-0.12) = 0$ and $F_v(0.12) = 1$. The distributions of the disturbance w_k and the initial estimation error ϵ_0 are defined analogously in terms of truncated normal distributions and bounds:

$$w_k' \sim \mathcal{N}(0, I/24^2),\quad \|w_k\|_\infty \le 0.12$$
$$\epsilon_0' \sim \mathcal{N}(0, 1/24^2),\quad \|\epsilon_0\| \le 0.12.$$

In accordance with the assumptions of section 1, the sequences $\{w_k, k = 0, 1, \dots\}$ and $\{v_k, k = 0, 1, \dots\}$ are i.i.d.

The feedback gain in (4b) was chosen as $K = [-0.847\ \ 1.00]$, which is the unconstrained LQ-optimal gain for the given plant model and cost. The observer gain was chosen to be the steady state Kalman Filter gain, $L = [2.11\ \ 0.21]^T$, which gives the spectral radius of $A - LC$ as 0.51.

The mode 1 prediction horizon N and the horizon ν over which constraints are handled explicitly were chosen (by

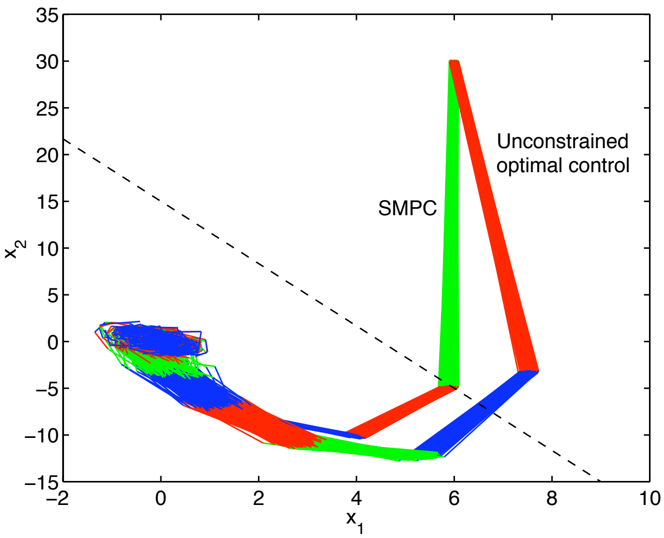

Fig. 1. Closed loop responses under SMPC and optimal unconstrained control laws, and the state constraint (dashed line).

considering the transient response when constraints are absent) as $N = 6$ and $\nu = 13$. The sequence $\{\hat{\gamma}_{0|k},\ k = 0, 1, \ldots, \nu{-}1\}$ was computed by discrete convolution based on an equi-spaced grid with spacing fixed at 10^{-4}. The values of $\beta_{i|k}$ were bounded using (18). This gives the condition of (21) as $h \geq \bar{\beta} = 3.94$, which is clearly satisfied by the given constraints. For the chosen N and ν, the terminal sets $\mathcal{S}_{\nu|k}$ and $\mathcal{S}_{\nu}$ are given by (23) and (24) with $N_k^* = 7$ for all k.

For 10^4 realizations of the disturbance and noise sequences, the response of the SMPC algorithm of section 4 was simulated and compared with that of the optimal unconstrained feedback law. In each simulation, the initial state was defined as $x_0 = \hat{x}_0 + \epsilon_0$, with the initial state estimate fixed: $\hat{x}_0 = (6, 30)$, and with the initial estimation error ϵ_0 obtained by sampling the distribution $\mathfrak{D}_0$.

The evolution of the state sequences $\{x_k, k = 0, 1, \ \ldots\}$ for 200 realizations of the disturbance and noise sequences is shown in Figure 1, which also shows the state constraint. For SMPC the observed probability of violating the constraint was 19.8% (at $k = 1$), which is close to the specified value of 20%, whereas violation rate was 100% at $k = 1$ under the unconstrained optimal feedback law. The cumulative costs:

$$\sum_{k=0}^{12} (x_k^T Q x_k + R u_k^2)$$

for closed loop operation averaged over 10^4 simulations were 1603.0 and 1580.9 under SMPC and unconstrained optimal control respectively. The small difference between these costs (1.38%) implies that the constraints can be satisfied with a specified probability at a very small cost.

6. CONCLUSION

This paper discusses and proposes a solution for a stochastic MPC problem for systems with random additive disturbances and soft constraints, and for which the states are not available for direct measurement. The proposed algorithm carries a guarantee of recursive feasibility and ensures a form of the quadratic stability and convergence of the state to a neighbourhood of the origin. Simulation results show that it achieves near-optimal performance while satisfying the soft constraints non-conservatively, in the sense that the observed probability of constraint satisfaction in closed loop operation is the same as the specified value.

Acknowledgements

Qifeng Cheng is supported by the K.C. Wong foundation.

REFERENCES

Cannon, M., Kouvaritakis, B., Raković, S., and Cheng, Q. (2010). Stochastic tubes in model predictive control with probabilistic constraints. *IEEE Trans. Automatic Control*. Accepted for publication.

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

Gilbert, E. and Tan, K. (1991). Linear systems with state and control constraints: The theory and practice of maximal admissible sets. *IEEE Trans. Automatic Control*, 36(9), 1008–1020.

Kouvaritakis, B., Cannon, M., Raković, S., and Cheng, Q. (2010). Explicit use of probabilistic distributions in linear predictive control. *Automatica*. Under review.

Kouvaritakis, B., Rossiter, J., and Schuurmans, J. (2000). Efficient robust predictive control. *IEEE Trans. Automatic Control*, 45(8), 1545–1549.

Langson, W., Chryssochoos, I., Raković, S., and Mayne, D. (2004). Robust model predictive control using tubes. *Automatica*, 40(1), 125–133.

Lee, Y. and Kouvaritakis, B. (2001). Receding horizon output feedback control for linear systems with input saturation. *IEE Proc. Control Theory Appl.*, 148(2), 109–115.

Løvaas, C., Seron, M., and Goodwin, G. (2008). Robust output-feedback model predictive control for systems with unstructured uncertainty. *Automatica*, 44(8), 1933–1943.

Mayne, D., Raković, S., Findeisen, R., and Allgöwer, F. (2006). Robust output feedback model predictive control of constrained linear systems. *Automatica*, 42(7), 1217–1222.

Mayne, D., Raković, S., Findeisen, R., and Allgöwer, F. (2009). Robust output feedback model predictive control of constrained linear systems: Time varying case. *Automatica*, 45(9), 2082–2087.

Mayne, D., Rawlings, J., Rao, C., and Scokaert, P. (2000). Constrained model predictive control: Stability and optimality. *Automatica*, 36(6), 789–814.

Mayne, D., Seron, M., and Raković, S. (2005). Robust model predictive control of constrained linear systems with bounded disturbances. *Automatica*, 41(2), 219–224.

Wan, Z. and Kothare, M. (2002). Robust output feedback model predictive control using off-line linear matrix inequalities. *J. Process Control*, 12(7), 763–774.

Yan, J. and Bitmead, R. (2005). Incorporating state estimation into model predictive control and its application to network traffic control. *Automatica*, 41, 595–604.