Proceedings of the 17th World Congress
The International Federation of Automatic Control
Seoul, Korea, July 6-11, 2008

Probabilistic constrained MPC for systems with multiplicative and additive stochastic uncertainty

Mark Cannon [*] **Basil Kouvaritakis** [*] **Xingjian Wu** [*]

[*] *Department of Engineering Science, University of Oxford, Parks Road, Oxford, OX1 3PJ, UK, mark.cannon@eng.ox.ac.uk*

Abstract: The paper develops a receding horizon control strategy to guarantee closed loop convergence and feasibility in respect of soft constraints. Earlier work (Cannon et al., 2007) presented results addressing closed loop stability in the case of multiplicative uncertainty only. The present paper extends these results to the more general case of additive and multiplicative uncertainty and proposes a method of handling probabilistic constraints. The results are illustrated by a design study considering control of a wind turbine in order to maximize power capture subject to constraints on fatigue damage.

Keywords: stochastic control, optimization, constrained control, fatigue life.

1. INTRODUCTION

Model predictive control (MPC) strategies are uniquely suited to using available information on the current state to optimize predicted behaviour in the presence of constraints, thereby providing a practicable solution to a closed-loop infinite dimensional optimization problem (Scokaert and Rawlings, 1998). Most applications however, apart from being constrained, are also subject to uncertainty. This could be handled by robust MPC techniques (Kothare et al., 1996; Mayne et al., 2000) which assume an uncertainty description that consists of bounds on the unknown parameters. However, such an approach is conservative because it fails to take account of information that is often available about the distribution of uncertainty. The difficulty here is that a systematic, non-conservative, efficient means for propagating the effects of uncertainty over a prediction horizon is needed; this remains largely an open question despite several important contributions (Batina et al., 2002; van Hessem et al., 2001).

A computationally convenient approach was proposed in Cannon et al. (2007), based on an autonomous description of the dynamics governing the evolution of future input and state predictions (Kouvaritakis et al., 2000). The paper uses the concept of probabilistic invariance to propose a strategy for handling soft constraints in the case of multiplicative uncertainty. The purpose of the current paper is to extend the approach to the more challenging problem of systems subject to both additive and multiplicative uncertainty. The presence of additive uncertainty implies that it is not possible to design control laws on the basis of mean square stability criteria, under which the convergence of state predictions would be ensured (with probability 1). The paper therefore determines the non-zero asymptotic limit of predicted behaviour and redefines the control objective based on deviations away from this limit. The paper then uses probabilistic invariance to construct an algorithm that provides desirable closed loop properties: asymptotic convergence of the state variance and feasibility of particular types of soft constraints.

Motivated by problems involving fatigue constraints, the soft constraints considered here take the form of limits on the expected number of samples at which a generalized system output exceeds specified bounds over a given horizon. A simplified analysis method is presented to demonstrate the effectiveness of probabilistic invariance in converting these constraints into probabilistic state constraints. The results of the paper are illustrated by a design example based on a simulated wind turbine control problem. The aim is to maximize profit by maximizing power capture while respecting constraints on high cycle fatigue damage (due to wind speed fluctuations) in order to prevent fatigue failure of the turbine blades over the required lifetime of the turbine.

2. PROBLEM FORMULATION

Consider the uncertain linear system described by

$$x_{k+1} = A_k x_k + B_k u_k + d_k, \quad x \in \mathbb{R}^{n_x}, \quad u \in \mathbb{R}^{n_u}. \qquad (1)$$

Let $\theta_k \in \mathbb{R}^m$ denote the vector of elements of A_k, B_k, d_k that are subject to uncertainty, and assume that $\{\theta_k,\ k = 0, 1, \ldots\}$ is a temporally independent and identically distributed (i.i.d.) sequence of Gaussian random variables: $\theta_k \sim \mathcal{N}(\bar{\theta}, \Sigma_\theta)$. Then $\theta_k = \bar{\theta} + U\Lambda^{1/2} q_k$, for i.i.d. $q_k \sim \mathcal{N}(0, I)$ (where U, Λ denote eigenvector and eigenvalue matrices of the covariance matrix Σ_θ), so the uncertainty description can be formulated as

$$[A_k\ B_k\ d_k] = \begin{bmatrix}\bar{A}\ \bar{B}\ 0\end{bmatrix} + \sum_{j=1}^{m} \begin{bmatrix}\tilde{A}_j\ \tilde{B}_j\ \tilde{g}_j\end{bmatrix} q_{k,j} \qquad (2)$$

with $q_k = [q_{k,1} \quad \cdots \quad q_{k,m}]^T \sim \mathcal{N}(0, I)$.

We define predicted future control sequences using a dual mode prediction paradigm (Mayne et al., 2000), according to which the control inputs over the first N steps are free variables, and a prescribed state feedback law is assumed over the subsequent infinite prediction horizon. Therefore the future control sequence predicted at time k, $\{u_{k+i},\ i = 0, 1, \ldots\}$, can be formulated as (e.g. Rossiter et al., 1998):

978-1-1234-7890-2/08/$20.00 © 2008 IFAC 10.3182/20080706-5-KR-1001.1048

$$u_{k+i} = \begin{cases} Kx_{k+i} + c_{i|k} & i = 0, \ldots, N-1 \\ Kx_{k+i} & i = N, N+1, \ldots \end{cases} \quad (3)$$

The dependence of u_{k+i} on the predicted state trajectory x_{k+i} (which is necessary due to the feedback law used in predictions for $i \geq N$) coupled with the Gaussian assumption of (2) imply that input and state constraints must be soft. This means that constraint violations may occur at any given time, provided that the likelihood of constraint violation does not exceed a specified bound. To make this more precise, define ψ_k as a system output, which may be subject to additive and multiplicative uncertainty:

$$\begin{aligned} \psi_k &= C_k x_k + D_k u_k + \eta_k, \quad \psi_k \in \mathbb{R}^n_\psi \\ [C_k \; D_k \; \eta_k] &= [\bar{C} \; \bar{D} \; 0] + \sum_{j=1}^{m} [\tilde{C}_j \; \tilde{D}_j \; \tilde{\eta}_j] \, q_{k,j} \end{aligned} \quad (4)$$

with $q_k = [q_{k,1} \;\; \cdots \;\; q_{k,m}]^T \sim \mathcal{N}(0, I)$.

We consider the constraint that the expected number of samples at which ψ_k lies outside a desired interval $I_\psi = [\psi_L, \psi_U]$ over a future horizon N_c should not exceed a bound $N_{\max}$:

$$\frac{1}{N_c} \sum_{i=0}^{N_c - 1} \Pr\{\psi_{k+i} \notin I_\psi\} \leq N_{\max}/N_c. \quad (5)$$

This statement can be translated into probabilistic constraints on the model state (as discussed in section 4), and hence into constraints invoked in the online MPC optimization (discussed in section 6). Within this framework soft input constraints are a special case of (4),(5) with $C_k = 0$, $D_k = I$, $\eta_k = 0$.

The use of feedback and feedforward actions prescribed by (3) leads to an autonomous description of the prediction dynamics (Kouvaritakis et al., 2000). Thus predictions at time k are generated by

$$\begin{aligned} z_{i+1|k} &= \Psi_{k+i} z_{i|k} + \delta_{k+i}, \quad i = 0, 1, \ldots \\ z_{0|k} &= \begin{bmatrix} x_k \\ f_k \end{bmatrix}, \quad f_k = \begin{bmatrix} c_{0|k} \\ \vdots \\ c_{N-1|k} \end{bmatrix} \\ [\Psi_k \; \delta_k] &= [\bar{\Psi} \; 0] + \sum_{j=1}^{m} [\tilde{\Psi}_j \; \tilde{\gamma}_j] \, q_{k,j} \\ \bar{\Psi} &= \begin{bmatrix} \bar{\Phi} & \bar{B}E \\ 0 & M \end{bmatrix}, \quad \tilde{\Psi}_j = \begin{bmatrix} \tilde{\Phi}_j & \tilde{B}_j E \\ 0 & 0 \end{bmatrix}, \quad \tilde{\gamma}_j = \begin{bmatrix} \tilde{g}_j \\ 0 \end{bmatrix} \end{aligned} \quad (6)$$

where the plant state x_k is assumed measurable at time k, $\bar{\Phi} = \bar{A} + \bar{B}K$, $\tilde{\Phi}_j = \tilde{A}_j + \tilde{B}_j K$, and

$$M = \begin{bmatrix} 0 & I & 0 & \cdots & 0 \\ 0 & 0 & I & \cdots & 0 \\ \vdots & \vdots & & & \vdots \\ 0 & 0 & 0 & \cdots & 0 \end{bmatrix}, \quad E = [I \;\; 0 \;\; \cdots \;\; 0]$$

with I denoting the identity matrix in $\mathbb{R}^{n_u \times n_u}$. In this formulation, the plant state x_{k+i} predicted at time k is then given by $x_{k+i} = \Gamma^T z_i$, where $\Gamma = [I \;\; 0]^T$.

The advantage of (6) over the conventional formulation of prediction dynamics is that constraints on a predicted output sequence at time k: $\{\psi_{k+i}, \; i = 0, 1, \ldots\}$ can be invoked conveniently through constraints on the initial prediction system state $z_{0|k}$. Since (6) is autonomous, the implied constraint set for $z_{0|k}$ can be computed using one-step-ahead invariance considerations. This provides significant computational advantages in the deterministic case, but in the stochastic case it is of vital importance since it removes the need to propagate the effects of uncertainty over the prediction horizon. Given that both the predicted states and the model parameters are stochastic, such propagation would require the distributions of future states to be computed numerically (see Batina et al., 2002; Li et al., 2002), implying a prohibitive computational load.

3. RECEDING HORIZON PERFORMANCE OBJECTIVE

Let $\mathbb{E}_k(X)$ denote the expected value of a random variable X, conditional on information available at time k, namely the state measurement x_k. This section considers the predictions made at a single instant k, and to simplify notation we therefore write $z_i = z_{i|k}$. Under the assumption that (6) is mean square stable, it can be shown (see Kushner, 1971) that the covariance matrix $\mathbb{E}_k(z_i z_i^T)$ converges to a finite limit as $i \to \infty$ along the predicted trajectories of (6), provided that Ψ_{k+i} and δ_{k+i} are independent. For completeness we give the extension of this result to the case that Ψ_{k+i} and δ_{k+i} are not independent.

Lemma 1. The sequence $\{z_i, \; i = 0, 1, \ldots\}$ generated by (6) satisfies $\lim_{i\to\infty} \mathbb{E}_k(z_i) = 0$ and $\lim_{i\to\infty} \mathbb{E}_k(z_i z_i^T) = \Theta$, where Θ is the solution of the Lyapunov equation

$$\Theta - \bar{\Psi}\Theta\bar{\Psi}^T - \sum_{j=1}^{m} \tilde{\Psi}_j \Theta \tilde{\Psi}_j^T = \sum_{j=1}^{m} \tilde{\gamma}_j \tilde{\gamma}_j^T \quad (7)$$

if and only if there exists $P > 0$ satisfying

$$P - \bar{\Psi}^T P \bar{\Psi} - \sum_{j=1}^{m} \tilde{\Psi}_j^T P \tilde{\Psi}_j > 0. \quad (8)$$

Proof: Given the linearity of (6), the sequence $\{z_i\}$ is the sum of the sequences $\{\zeta_i\}$ and $\{\xi_i\}$ generated by the following two systems:

$$\zeta_{i+1} = \Psi_{k+i}\zeta_i, \qquad \zeta_0 = z_0 \quad (9a)$$

$$\xi_{i+1} = \Psi_{k+i}\xi_i + \delta_{k+i}, \qquad \xi_0 = 0 \quad (9b)$$

Condition (8) is necessary and sufficient (see e.g. Boyd et al., 1994) for mean square stability of (9a) and therefore ensures that $\mathbb{E}_k(\zeta_i \zeta_i^T) \to 0$, and hence $\zeta_i \to 0$ almost surely, as $i \to \infty$. From (9b) we have $\mathbb{E}_k(\xi_i) = 0$ for all i, and it follows that $\mathbb{E}_k(z_i) \to 0$ as $i \to \infty$. Since ξ_i is independent of Ψ_{k+i}, (9b) also gives

$$\begin{aligned} \mathbb{E}_k(\xi_{i+1}\xi_{i+1}^T) &= \mathbb{E}_k\{(\Psi_{k+i}\xi_i + \delta_{k+i})(\Psi_{k+i}\xi_i + \delta_{k+i})^T\} \\ &= \mathbb{E}_k(\Psi_{k+i}\xi_i\xi_i^T\Psi_{k+i}^T) + \mathbb{E}_k(\delta_{k+i}\delta_{k+i}^T) \\ &= \bar{\Psi}\mathbb{E}_k(\xi_i\xi_i^T)\bar{\Psi}^T + \sum_{j=1}^{m} \tilde{\Psi}_j \mathbb{E}_k(\xi_i\xi_i^T)\tilde{\Psi}_j^T + \sum_{j=1}^{m} \tilde{\gamma}_j\tilde{\gamma}_j^T. \end{aligned} \quad (10)$$

Let $\hat{\Theta}_i = \mathbb{E}_k(\xi_i\xi_i^T) - \Theta$, then from (7) and (10) we have

$$\hat{\Theta}_{i+1} = \bar{\Psi}\hat{\Theta}_i\bar{\Psi}^T + \sum_{j=1}^{m} \tilde{\Psi}_j \hat{\Theta}_i \tilde{\Psi}_j^T.$$

From the mean square stability condition of (8), it follows that $\hat{\Theta}_i \to 0$, so that $\mathbb{E}_k(\xi_i\xi_i^T) \to \Theta$ as $i \to \infty$. Finally, note that $\mathbb{E}_k(z_i z_i^T) \to \mathbb{E}_k(\xi_i\xi_i^T)$ since $\zeta_i \to 0$ as $i \to \infty$. □

For the case of no additive uncertainty, i.e. if $\tilde{\gamma}_j = 0$ for $j = 1, \ldots, m$ in (6), the existence of P satisfying (8) implies that $\lim_{i\to\infty} \mathbb{E}(z_i z_i^T) = 0$. In this case therefore, the predicted cost

$$\begin{aligned} J_k &= \sum_{i=0}^{\infty} \mathbb{E}_k(x_{k+i}^T Q x_{k+i} + u_{k+i}^T R u_{k+i}) \\ &= \sum_{i=0}^{\infty} \mathbb{E}_k(z_i^T \tilde{Q} z_i), \quad \tilde{Q} = \begin{bmatrix} Q + K^T R K & K^T R E \\ E^T R K & E^T R E \end{bmatrix} \end{aligned} \tag{11}$$

is well-defined, and hence is a suitable candidate for a predicted performance cost to be minimized online by a receding horizon control law. However in the case of persistent non-zero additive uncertainty, it follows from Lemma 1 that the stage cost of (11) converges to a nonzero limit along trajectories of (6):

$$\lim_{i\to\infty} \mathbb{E}_k(z_i^T \tilde{Q} z_i) = \mathrm{tr}(\Theta \tilde{Q})$$

(where $\mathrm{tr}(M)$ denotes the trace of M). Therefore the cost (11) is infinite in this case, and to obtain a finite cost, the predicted cost defined in (11) must be modified:

$$J_k = \sum_{i=0}^{\infty} \mathbb{E}_k(L_i), \quad L_i = z_i^T \tilde{Q} z_i - \mathrm{tr}(\Theta \tilde{Q}). \tag{12}$$

The following result shows that the cost (12) can be evaluated as a quadratic function of the initial state of the prediction dynamics (6).

Theorem 2. The cost (12), evaluated along trajectories of (6), is given by

$$J_k = \begin{bmatrix} z_0 \\ 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} z_0 \\ 1 \end{bmatrix}, \quad \tilde{P} = \begin{bmatrix} P_z & P_{z1} \\ P_{1z} & P_1 \end{bmatrix}$$

where P_z, $P_{1z} = P_{z1}^T$, and P_1 are defined uniquely by

$$P_z - \bar{\Psi}^T P_z \bar{\Psi} - \sum_{j=1}^{m} \tilde{\Psi}_j^T P_z \tilde{\Psi}_j = \tilde{Q} \tag{13a}$$

$$P_{1z} = \sum_{j=1}^{m} \tilde{\gamma}_j^T P_z \tilde{\Psi}_j (I - \bar{\Psi})^{-1} \tag{13b}$$

$$P_1 = -\,\mathrm{tr}(\Theta P_z). \tag{13c}$$

Proof: Define $V_i = z_i^T P_z z_i + z_i^T P_{z1} + P_{1z} z_i + P_1$, then from $z_{i+1} = \Psi_{k+i} z_i + \delta_{k+i}$ we have

$$\begin{aligned} &\mathbb{E}_k(V_i) - \mathbb{E}_k(V_{i+1}) = \mathbb{E}_k\big(z_i^T \big[P_z - \mathbb{E}(\Psi_{k+i}^T P_z \Psi_{k+i})\big] z_i\big) \\ &+2\big[P_{1z}(I - \bar{\Psi}) - \mathbb{E}(\delta_{k+i}^T P_z \Psi_{k+i})\big]\mathbb{E}_k(z_i) - \mathbb{E}(\delta_{k+i}^T P_z \delta_{k+i}). \end{aligned} \tag{14}$$

But (13a) gives the first term of the RHS as

$$\mathbb{E}_k\big(z_i^T \big[P_z - \mathbb{E}(\Psi_{k+i}^T P_z \Psi_{k+i})\big] z_i\big) = \mathbb{E}_k(z_i^T \tilde{Q} z_i). \tag{15}$$

Post-multiplying (7) by P_z and extracting the trace gives

$$\mathrm{tr}(\Theta P_z - \bar{\Psi}\Theta\bar{\Psi}^T P_z - \sum_{j=1}^{m} \tilde{\Psi}_j \Theta \tilde{\Psi}_j^T P_z) =$$

$$\mathrm{tr}(\Theta[P_z - \bar{\Psi}^T P_z \bar{\Psi} - \sum_{j=1}^{m} \tilde{\Psi}_j^T P_z \tilde{\Psi}_j]) = \sum_{j=1}^{m} \mathrm{tr}(\tilde{\gamma}_j \tilde{\gamma}_j^T P_z)$$

and (13a) therefore implies

$$\mathrm{tr}(\Theta \tilde{Q}) = \sum_{j=1}^{m} \tilde{\gamma}_j^T P_z \tilde{\gamma}_j = \mathbb{E}(\delta_{k+i}^T P_z \delta_{k+i}). \tag{16}$$

Substituting (15) and (16) into (14), and using (13b) gives

$$\mathbb{E}_k(V_i) - \mathbb{E}_k(V_{i+1}) = \mathbb{E}_k(z_i^T \tilde{Q} z_i) - \mathrm{tr}(\Theta \tilde{Q}),$$

and by summing this recursion over all $i \geq 0$ we obtain

$$V_0 - \lim_{i\to\infty} \mathbb{E}_k(V_i) = \sum_{i=0}^{\infty} \mathbb{E}_k(L_i) = J_k.$$

The proof is completed by showing that $\mathbb{E}_k(V_i) \to 0$ as $i \to \infty$. This follows from the definition of V_i and (13c), which give

$$\mathbb{E}_k(V_i) = \mathbb{E}_k(z_i^T P_z z_i) + 2P_{1z}\mathbb{E}_k(z_i) - \mathrm{tr}(\Theta P_z)$$

and therefore

$$\begin{aligned} \lim_{i\to\infty} \mathbb{E}_k(V_i) &= \lim_{i\to\infty} \mathbb{E}_k(z_i^T P_z z_i) - \mathrm{tr}(\Theta P_z) \\ &= \lim_{i\to\infty} \mathrm{tr}\big[\mathbb{E}_k(z_i z_i^T) P_z\big] - \mathrm{tr}(\Theta P_z) = 0 \end{aligned}$$

where $\lim_{i\to\infty} \mathbb{E}_k(z_i) = 0$ and $\lim_{i\to\infty} \mathbb{E}_k(z_i z_i^T) = \Theta$ (from Lemma 1) have been used. □

4. A FRAMEWORK FOR HANDLING SOFT CONSTRAINTS

This section describes a method of analysis that enables the conversion of soft constraints (5) into probabilistic constraints on the state of (1). Let $\mathcal{E}_1 \subset \mathcal{E}_2 \subset \mathbb{R}^{n_x}$ and assume that x_k can lie in either $S_1 = \mathcal{E}_1$ or $S_2 = \mathcal{E}_2 - \mathcal{E}_1$. This scenario contravenes the assumption that the uncertainty in (2) has infinite support, but it is based on the assumption that $\mathcal{E}_2$ is defined so that the probability of $x_k \notin \mathcal{E}_2$ is negligible. The analysis could be made less conservative (and more realistic) by considering a sequence of nested sets: $\mathcal{E}_1 \subset \cdots \subset \mathcal{E}_r$, but $r = 2$ is used here to simplify presentation.

Define the conditional probabilities

$$\Pr(\psi_k \notin I_\psi \mid x_k \in S_j) = p_j, \quad j = 1, 2. \tag{17}$$

Under the assumption that p_1 is small, so that S_1 is the *safe* region of state space, it is convenient (though possibly conservative), to assume that S_2 is *unsafe* and set $p_2 = 1$. Define also the matrix of transition probabilities

$$\Pi = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix} \tag{18}$$

where p_{ij} is the probability that the online algorithm steers the state in one step from from S_j to S_i. Then over i steps we have

$$\begin{bmatrix} \Pr(x_{k+i} \in S_1) \\ \Pr(x_{k+i} \in S_2) \end{bmatrix} = \Pi^i \begin{bmatrix} \Pr(x_k \in S_1) \\ \Pr(x_k \in S_2) \end{bmatrix}$$

so that the probability of a constraint violation at time $k + i$ is given by

$$\Pr(\psi_{k+i} \notin I_\psi) = [p_1 \;\, p_2]\, \Pi^i \begin{bmatrix} \Pr(x_k \in S_1) \\ \Pr(x_k \in S_2) \end{bmatrix}$$

Because of the special structure of Π (see Kushner, 1971), its eigenvalue/vector decomposition is given by

$$\Pi = [w_1 \;\, w_2] \begin{bmatrix} 1 & 0 \\ 0 & \lambda_2 \end{bmatrix} \begin{bmatrix} v_1^T \\ v_2^T \end{bmatrix}, \quad 0 \leq \lambda_2 < 1.$$

Therefore, as $i \to \infty$, the rate at which constraint violations accumulate given $x_k \in S_j$ tends to

$$R_j = [p_1 \;\, p_2]\, w_1 v_1^T e_j, \quad j = 1, 2 \tag{19}$$

where e_1, e_2 denote the first and second columns of the 2×2 identity matrix. If R_1 and R_2 are less than the limit

$N_{\max}/N_c$ of (5), then it follows that there exists finite i^* such that, for all $i \geq i^*$, the total expected number of constraint violations will be less than $i^* N_{\max}/N_c$. Provided i^* does not exceed the horizon N_c, it then follows that the probabilistic formulation (17),(18) ensures that the soft constraints (5) on ψ are satisfied.

5. PROBABILISTIC INVARIANCE AND PROBABILISTIC CONSTRAINTS

In the scenario discussed in section 4, the satisfaction of soft constraints on ψ is dependent on the probability p_i of constraint violation given that the state lies in the set $\mathcal{E}_i$ and on the transition probabilities p_{ij} between these sets. This section proposes a procedure for designing $\mathcal{E}_1$ and assigning probabilities to p_1 and p_{11}. This is done below using the concept of probabilistic invariance, which is defined as follows.

Definition 3. (Cannon et al., 2007). A set $S \subset \mathbb{R}^{n_x}$ is invariant with probability p (i.w.p. p) under a given control law if $x_{k+1} \in S$ with probability p whenever $x_k \in S$.

The approach is based on ellipsoidal sets, $\mathcal{E} \subset \mathbb{R}^{n_x + Nn_u}$ and $\mathcal{E}_x \subset \mathbb{R}^{n_x}$ defined:

$$\mathcal{E} = \{z : z^T \hat{P} z \leq 1\}$$
$$\mathcal{E}_x = \{x : x^T \hat{P}_x x \leq 1\}, \quad \hat{P}_x = (\Gamma^T \hat{P}^{-1} \Gamma)^{-1}$$

so that $\mathcal{E}_x$ is the projection of $\mathcal{E}$ onto the x-subspace (i.e. f exists such that $z = [x^T \ \ f^T]^T \in \mathcal{E}$ whenever $x \in \mathcal{E}_x$). Let $\mathcal{Q}$ denote a set that contains the vector of uncertain parameters in (2) with a specified probability p:

$$\Pr\{q_k \in \mathcal{Q}\} \geq p. \tag{20}$$

Since $\|q_k\|_2^2$ has a chi-square distribution with m degrees of freedom, a set with this property is the hypersphere $\{q : \|q\|_2 \leq r\}$, where $\Pr(\chi^2(m) \leq r^2) = p$. Earlier work (Cannon et al., 2007) used ellipsoidal confidence regions derived from this hypersphere to compute i.w.p. sets, but to accommodate the additive uncertainty in (1), we assume here that $\mathcal{Q}$ is polytopic with vertices $q^{(i)}$, $i = 1, \ldots, \nu$. Thus any polytope containing the hypersphere $\{q : \|q\|_2 \leq r\}$ provides a convenient (possibly conservative) choice for $\mathcal{Q}$. The following Lemma gives conditions for probabilistic invariance of $\mathcal{E}_x$.

Lemma 4. $\mathcal{E}_x$ is i.w.p. p under (3) for any f_k such that $z_{0|k} \in \mathcal{E}$ if there exists a scalar $\lambda \in [0, 1]$ satisfying

$$\begin{bmatrix} \hat{P}_x^{-1} & \Gamma^T \Psi(q^{(i)}) \hat{P}^{-1} & \Gamma^T \gamma(q^{(i)}) \\ \hat{P}^{-1} \Psi(q^{(i)})^T \Gamma & \lambda \hat{P}^{-1} & 0 \\ \gamma(q^{(i)})^T \Gamma & 0 & 1 - \lambda \end{bmatrix} \geq 0 \tag{21}$$

for $i = 1, \ldots, \nu$, where $\Psi(q_k) = \bar{\Psi} + \sum_{j=1}^m \tilde{\Psi}_j q_{k,j}$ and $\gamma(q_k) = \sum_{j=1}^m \tilde{\gamma}_j q_{k,j}$.

Proof: From (20) it follows that $\Pr(x_{k+1}^T \hat{P}_x x_{k+1} \leq 1) \geq p$ if

$$x_{k+1}^T \hat{P}_x x_{k+1} \leq 1 \quad \forall z_{0|k} \in \mathcal{E}, \ \ \forall q_k \in \mathcal{Q} \tag{22}$$

where, under (3), x_{k+1} is given by $x_{k+1} = \Gamma^T \Psi(q_k) z_{0|k} + \Gamma^T \gamma(q_k)$. By the $\mathcal{S}$-procedure (22) is equivalent to the existence of $\lambda \geq 0$ satisfying

$$1 - (\Psi(q)z + \gamma(q))^T \Gamma \hat{P}_x \Gamma^T (\Psi(q)z + \gamma(q)) \geq \lambda(1 - z^T \hat{P} z)$$

for all z and all $q \in \mathcal{Q}$, or equivalently

$$\begin{bmatrix} \lambda \hat{P} & 0 \\ 0 & 1 - \lambda \end{bmatrix} - \begin{bmatrix} \Psi(q)^T \\ \gamma(q)^T \end{bmatrix} \Gamma \hat{P}_x \Gamma^T \left[\Psi(q) \ \ \gamma(q)\right] \geq 0, \quad \forall q \in \mathcal{Q}.$$

Using Schur complements this can be expressed as a LMI in q, which, when invoked for all $q \in \mathcal{Q}$ is equivalent to (21) for some $\lambda \in [0, 1]$. □

Additional constraints on $\hat{P}$ are needed in order to constrain the conditional probability that ψ_k lies outside the desired interval I_ψ given that $z_{0|k}$ lies in $\mathcal{E}$. Re-writing (4) in the form:

$$\psi_k = \hat{C}(q_k) z_{0|k} + \eta(q_k)$$
$$\hat{C}(q_k) = \bar{C}\Gamma^T + \bar{D}\hat{K} + \sum_{j=1}^m (\tilde{C}_j \Gamma^T + \tilde{D}_k \hat{K}) q_{k,j}$$

with $\hat{K} = [K \ \ E]$, the following result is based on the confidence region $\mathcal{Q}$.

Lemma 5. $\Pr(\psi_k \notin I_\psi \mid z_{0|k} \in \mathcal{E}) \leq 1 - p$ if

$$\psi_L \leq \eta(q^{(i)}) \leq \psi_U \tag{23}$$

for $i = 1, \ldots, \nu$ and

$$\left[\hat{C}(q^{(i)}) \hat{P}^{-1} \hat{C}(q^{(i)})^T\right]_{jj} \leq \left[\psi_U - \eta(q^{(i)})\right]_j^2 \tag{24a}$$
$$\left[\hat{C}(q^{(i)}) \hat{P}^{-1} \hat{C}(q^{(i)})^T\right]_{jj} \leq \left[\eta(q^{(i)}) - \psi_L\right]_j^2 \tag{24b}$$

for $i = 1, \ldots, \nu$ and $j = 1, \ldots, n_\psi$, where $[\ ]_{ij}$ denotes element ij.

Proof: For given q, $\max_{z \in \mathcal{E}} [\hat{C}(q) z]_j = [\hat{C}(q) \hat{P}^{-1} \hat{C}(q)^T]_{jj}^{1/2}$, and it follows from (20) that $\Pr(\psi_k \in I_\psi) \geq p$ whenever $z_{0|k} \in \mathcal{E}$ if

$$\left[\hat{C}(q) \hat{P}^{-1} \hat{C}(q)^T\right]_{jj}^{1/2} \leq \left[\psi_U - \eta(q)\right]_j \tag{25a}$$
$$\left[\hat{C}(q) \hat{P}^{-1} \hat{C}(q)^T\right]_{jj}^{1/2} \leq \left[\eta(q) - \psi_L\right]_j \tag{25b}$$

for all $q \in \mathcal{Q}$ and $j = 1, \ldots, n_\psi$. Since (25a,b) are convex in q, the equivalent constraints (23),(24a,b) are obtained by invoking (25a,b) at each vertex of the polytope $\mathcal{Q}$. □

With $\mathcal{E}_1$ defined as $\mathcal{E}_x$ in the example outlined in section 4, the values of p_1 and p_{11} can specified using the constraints of lemmas 4 and 5. To maximize the *safe* region of state space, it is clearly desirable to maximize $\mathcal{E}_x$, which can be formulated as

$$\underset{\hat{P}^{-1}, \ \lambda \in [0,1]}{\text{maximize}} \ \ \det(\hat{P}_x^{-1}) \text{ subject to (21),(23),(24a,b)} \tag{26}$$

Remark 6. If λ is a constant, then the constraints in (26) are LMIs in $\hat{P}^{-1}$. Therefore $\mathcal{E}_x$ can be optimized by successively maximizing $\det(\hat{P}_x^{-1})$ over the variable $\hat{P}^{-1}$ subject to (21),(23),(24a,b), with the scalar λ fixed at a sequence of values in the interval $[0, 1]$.

Remark 7. In the case of input constraints: $u_L \leq u_k \leq u_U$, the constraints of lemma 5 reduce to $u_L \leq 0 \leq u_U$ and

$$\left[\hat{K}^T \hat{P}^{-1} \hat{K}\right]_{jj} \leq \left[u_U\right]_j^2, \quad \left[\hat{K}^T \hat{P}^{-1} \hat{K}\right]_{jj} \leq \left[u_L\right]_j^2 \tag{27}$$

for $j = 1, \ldots, n_u$. In the example of section 4 with $\mathcal{E}_1 = \mathcal{E}_x$, this implies $p_1 = 0$.

6. RECEDING HORIZON CONTROL

If the plant state is maintained inside the set $\mathcal{E}_x$ with probability p_{11} and is returned to $\mathcal{E}_x$ with probability greater than or equal to p_{12}, then by the arguments of section 4, both the predicted and closed loop responses are guaranteed to satisfy the constraints (5). The aim of the online MPC is therefore two-fold: (a) minimize the cost (12) subject to $z_{0|k} \in \mathcal{E}$ whenever $x_k \in \mathcal{E}_x$; or (b) return the state to $\mathcal{E}_x$ as quickly as possible whenever $x_k \notin \mathcal{E}_x$. The latter is achieved by driving the expected value of x_{k+1} as close to (or as far inside) $\mathcal{E}_x$ as possible. This strategy is an indirect but computationally convenient means of increasing the value of p_{12}. The algorithm can be stated as follows.

Algorithm 1. At times $k = 0, 1, \ldots$:
1. If $x_k \in \mathcal{E}_x$, compute

$$f_k^* = \arg\min_{f_k} \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix} \quad \text{subject to} \quad z_{0|k}^T \hat{P} z_{0|k} \leq 1. \tag{28}$$

2. If $x_k \notin \mathcal{E}_x$, compute

$$f_k^* = \arg\min_{f_k} z_{0|k}^T \bar{\Psi} \Gamma \hat{P}_x \Gamma^T \bar{\Psi} z_{0|k}$$

$$\text{subject to} \quad \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix} \leq \begin{bmatrix} x_k \\ M f_{k-1}^* \\ 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} x_k \\ M f_{k-1}^* \\ 1 \end{bmatrix}. \tag{29}$$

3. Implement $u_k = K x_k + E f_k^*$.

Both (28) and (29) require the minimization of a convex quadratic cost subject to a convex quadratic constraint, which can be solved efficiently (using e.g. the technique discussed in Kouvaritakis et al., 2002). The constraint in (29) is introduced in order to ensure that the time-average of the expected value of $z_{0|k}^T \tilde{Q} z_{0|k}$ converges to a finite limit, as shown by the following theorem.

Theorem 8. The closed loop response of (1) under algorithm 1 satisfies

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n} \mathbb{E}_0\left(\begin{bmatrix} x_k \\ f_k^* \end{bmatrix}^T \tilde{Q} \begin{bmatrix} x_k \\ f_k^* \end{bmatrix} \right) \leq \operatorname{tr}(\Theta \tilde{Q}). \tag{30}$$

Proof: Let J_k^* denote the value of J_k corresponding to the optimal f_k^* computed by algorithm 1 at time k. Then

$$\mathbb{E}_k(J_{k+1}^*) = \mathbb{E}_k(J_{k+1}^* | x_{k+1} \in \mathcal{E}_x) \Pr(x_{k+1} \in \mathcal{E}_x) + \mathbb{E}_k(J_{k+1}^* | x_{k+1} \notin \mathcal{E}_x) \Pr(x_{k+1} \notin \mathcal{E}_x)$$

where J_{k+1}^* necessarily satisfies

$$J_{k+1}^* \leq \begin{bmatrix} \Psi_k z_{0|k} + \delta_k \\ 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} \Psi_k z_{0|k} + \delta_k \\ 1 \end{bmatrix}$$

(on account of the objective in (28) if $x_{k+1} \in \mathcal{E}_x$ or because of the constraint in (29) if $x_{k+1} \notin \mathcal{E}_x$), and therefore

$$\mathbb{E}_k(J_{k+1}^*) \leq \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix}^T \mathbb{E}\left(\begin{bmatrix} \Psi_k & \delta_k \\ 0 & 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} \Psi_k & \delta_k \\ 0 & 1 \end{bmatrix} \right) \begin{bmatrix} z_{0|k} \\ 1 \end{bmatrix}. \tag{31}$$

However, from (13a,b,c) it follows that

$$\mathbb{E}\left(\begin{bmatrix} \Psi_k & \delta_k \\ 0 & 1 \end{bmatrix}^T \tilde{P} \begin{bmatrix} \Psi_k & \delta_k \\ 0 & 1 \end{bmatrix} \right) = \tilde{P} - \begin{bmatrix} \tilde{Q} & 0 \\ 0 & -\operatorname{tr}(\Theta \tilde{Q}) \end{bmatrix}$$

and (31) therefore implies that

$$J_k^* - \mathbb{E}_k(J_{k+1}^*) \geq z_{0|k}^T \tilde{Q} z_{0|k} - \operatorname{tr}(\Theta \tilde{Q}).$$

Recursion of this equation for $k = 0, 1, \ldots$ gives

$$\lim_{n\to\infty} \frac{1}{n} \big(J_0^* - \mathbb{E}_0(J_n^*) \big) \geq \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n} \mathbb{E}_0(z_{0|k}^T \tilde{Q} z_{0|k}) - \operatorname{tr}(\Theta \tilde{Q})$$

and, since J_n^* is lower bounded (because $P_z > 0$ in (13a)), it follows that

$$\lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n} \mathbb{E}_0(z_{0|k}^T \tilde{Q} z_{0|k}) \leq \operatorname{tr}(\Theta \tilde{Q})$$

which implies the time-average bound of (30). To complete the proof, note that the constraint in (28) is feasible whenever $x_k \in \mathcal{E}_x$ due to the definition of $\mathcal{E}_x$, and similarly the feasibility of $f_k = M f_{k-1}^*$ implies that the constraint in (29) is necessarily feasible. □

Corollary 9. If the probabilities p_{11}, p_{12} are such that, for the conditional constraint violation probability p_1, the expected rates R_1, R_2 of accumulation of constraint violations are within allowable limits, then the bound (5) will be satisfied in closed loop operation under algorithm 1.

Proof: This follows directly from the assumptions on R_1, R_2 and the arguments of section 4. □

The algorithm must be initialized by computing $\hat{P}$. A possible procedure for this is as follows: specify initial values p_{11}^0, p_{12}^0 for p_{11}, p_{12}. Then, from the bound $N_{\max}/N_c$ on the allowed rate of accumulation of constraint violations, the analysis of section 4 can be used to compute the minimum permissible value for p_1. Given p_{11}, p_{12} and p_1, the uncertainty set $\mathcal{Q}$ can be constructed and the constraints (21),(23),(24a,b) formulated, allowing $\hat{P}$ to be optimized by solving (26). Once $\hat{P}$ has been determined, the actual value of p_{12} can be computed (e.g. by Monte Carlo simulation); this must be greater than or equal to p_{12}^0 to ensure satisfaction of (5). If this is not the case, then $\hat{P}$ must be re-computed using reduced values for p_{11}^0, p_{12}^0. Note that the computation of $\hat{P}$ is performed offline.

Remark 10. If several desired intervals are specified for ψ, each with a bound on the expected number of violations, then the appropriate value for p_1 can be computed based on a weighted average rate of constraint violation. This situation is common when constraints on fatigue damage due to stress cycles of varying amplitudes are considered.

7. SIMULATION EXAMPLE

Consider the problem of maximizing the power capture of a variable pitch wind turbine while respecting limits on turbine blade fatigue damage caused by wind fluctuations. It is common practice to assume that the statistical properties of the wind remain constant over a period of order 10 minutes (Burton et al., 2001). Below rated average wind speed (but above cut-off wind speed), the control objective becomes that of maximizing efficiency, which can be achieved by regulating blade pitch angle about a given setpoint (determined by maximizing an appropriate function of wind speed, pitch angle, and blade angular velocity). This however is to be performed subject to constraints on the stress cycles experienced by the blades in order to achieve a specified fatigue life.

A simplified model of blade pitch rotation is given by

$$J \frac{d^2\beta}{dt^2} + c \frac{d\beta}{dt} = T_m - T_p \tag{32}$$

where β is the blade pitch angle, T_m is a torque applied by an actuator used to adjust β, and T_p is the pitching torque due to fluctuations in wind speed, which is a known function of wind speed and the blade's angle of attack, α. It should be noted that α is related (in a known manner) to wind speed and β. Therefore the model (32) is subject to additive stochastic uncertainty (due to the dependence of T_p on wind speed) and multiplicative uncertainty (due to the dependence of T_p on β), and furthermore these two sources of uncertainty are statistically dependent.

Blade fatigue damage depends on the resultant applied torque, so fatigue constraints are invoked on ψ defined by

$$\psi = T_m - T_p.$$

By considering variations about a given setpoint for β, a linear discrete model approximation was identified in the form of an ARMA model:

$$y_{k+1} = a_{k,1}y_k + a_{k,0}y_{k-1} + b_{k,1}u_k + b_{k,0}u_{k-1} + w_k \quad (33)$$

using data applied to a continuous-time model of the NACA 632-215(V) blade (Burton et al., 2001). Least squares estimates of $\theta = [a_1 \;\; a_0 \;\; b_1 \;\; b_0 \;\; w]^T$ were obtained from 1000 simulation runs, each with a given fixed wind speed. On the basis of these simulations, the mean $\bar{\theta}$ and covariance Σ_θ of the parameter vector θ were determined.

The model (33) can be written in the form (1), with

$$A_k = \begin{bmatrix} 0 & a_{k,2} \\ 1 & a_{k,1} \end{bmatrix}, \quad B_k = \begin{bmatrix} b_{k,2} \\ b_{k,1} \end{bmatrix}, \quad d_k = \begin{bmatrix} 0 \\ w_k \end{bmatrix}.$$

The identified parameters $(\bar{\theta}, \Sigma_\theta)$ indicate that B has negligible uncertainty. For a sampling interval of 1 second the corresponding uncertainty class is given by

$$[A_k \;\; d_k] = [\bar{A} \;\; 0] + \sum_{j=1}^{3} [\tilde{A}_j \;\;\; \tilde{g}_j]\, q_{k,j} \quad (34)$$

$$\bar{A} = \begin{bmatrix} 0 & -0.97 \\ 1 & 1.56 \end{bmatrix}, \; \bar{B} = \begin{bmatrix} -0.20 \\ -0.21 \end{bmatrix}, \; [\tilde{A}_1 \;\; \tilde{d}_1] = \begin{bmatrix} 0 & -0.09 & 0 \\ 0 & 0.13 & 0.02 \end{bmatrix}$$

$$[\tilde{A}_2 \;\; \tilde{d}_2] = \begin{bmatrix} 0 & 0.21 & 0 \\ 0 & -0.009 & -0.06 \end{bmatrix}, \; [\tilde{A}_3 \;\; \tilde{d}_3] = \begin{bmatrix} 0 & -0.06 & 0 \\ 0 & 0.02 & 0.05 \end{bmatrix}$$

The Gaussian assumption on q_k was validated by the Jarque-Bera test at the 5% level. A discrete-time linearized description of the output ψ_k was estimated using a similar approach. The uncertainty in D was found to be negligible, and the uncertainty class for $[C \;\eta]$ was formulated as

$$[C_k \;\; \eta_k] = [\bar{C} \;\; 0] + \sum_{j=1}^{2} [\tilde{C}_j \;\;\; \tilde{\eta}_j]\, q_{k,j} \quad (35)$$

$$\bar{C} = [0 \;\; 729]\,, \; \bar{D} = 959$$

$$[\tilde{C}_1 \;\; \eta_1] = [0 \;\; 300 \;\; 50]\,, \; [\tilde{C}_2 \;\; \eta_2] = [0 \;\; 50 \;\; 100]$$

The number of degrees of freedom in predictions (3) was chosen as $N = 4$, and $N_c = 4$ was also used as the horizon over which to invoke the upper bound $N_{\max}$ on the permissible number of constraint violations. Miner's rule was used to determine $N_{\max}/N_c$, assuming (for simplicity) a single threshold on the torque $T_m - T_p$. Accordingly, for $p^0_{11} = 0.9$, $p^0_{12} = 0.8$, $N_{\max}/N_c = 0.3$, the permissible value for p_1 was found to be 0.2. For these values, the optimization (26) resulted in

$$\hat{P}_x = \begin{bmatrix} 0.03 & 0.04 \\ 0.04 & 0.069 \end{bmatrix}$$

Closed loop simulations of algorithm 1, performed for an initial condition $x_0 = [-7.88 \;\; 7.31]^T$ (which is close to the boundary of $\mathcal{E}_x$), gave an average number of constraint violations of 3 over a horizon of 40 steps, while the maximum number of constraint violations on any one simulation run was 4. From these simulations, the actual value of p_{12} was found to be 0.85, which exceeds p^0_{12}, indicating that algorithm 1 satisfies the fatigue constraints.

To establish the efficacy of algorithm 1, closed loop simulations were performed for 1000 sequences of uncertainty realizations, and compared in terms of cost and constraint satisfaction with the mean square stabilizing linear feedback law $u_k = Kx_k$. Algorithm 1 gave an average closed loop cost of 257, whereas the average cost for $u_k = Kx_k$ was 325. Algorithm 1 achieves this improvement in performance by driving (during transients) the predictions hard against the limits of the soft constraints. Both algorithm 1 and $u_k = Kx_k$ on average resulted a total number of constraint violations within the specified limit over a 40-step horizon. This is to be expected since both control laws achieve acceptable rates of constraint violation in steady state. However the average numbers of constraint violations over n steps, for $0 < n \le 16$, indicate that $u_k = Kx_k$ exceeded the allowable limits during transients, whereas algorithm 1 gave average constraint violation rates less than $N_{\max}/N_c = 0.3$ for all $n \ge i^* = 4$.

REFERENCES

I. Batina, A.A. Stoorvogel, and S. Weiland. Optimal control of linear, stochastic systems with state and input constraints. In *Proc. 41st IEEE Conf. Decision and Control*, pages 1564–1569, 2002.

S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory.* SIAM, 1994.

T. Burton, D. Sharpe, N. Jenkins, and E. Bossanyi. *Wind Energy Handbook.* Wiley, 2001.

M. Cannon, B. Kouvaritakis, and X. Wu. Stochastic predictive control with probabilistic constraints. Submitted to Automatica, 2007.

M.V. Kothare, V. Balakrishnan, and M. Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32(10):1361–1379, 1996.

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

B. Kouvaritakis, M. Cannon, and J.A. Rossiter. Who needs QP for linear MPC anyway? *Automatica*, 38(5): 879–884, 2002.

H.J. Kushner. *Introduction to stochastic control.* Holt, Rinehart and Winston, 1971.

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

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

J.A. Rossiter, B. Kouvaritakis, and M.J. Rice. A numerically robust state space approach to stable predictive control strategies. *Automatica*, 34(1):65–73, 1998.

P.O.M. Scokaert and J.B. Rawlings. Constrained linear quadratic regulation. *IEEE Trans. Automatic Control*, 43(8):1163–1169, 1998.

D.H. van Hessem, C.W. Scherer, and O.H. Bosgra. LMI-based closed-loop economic optimization of stochastic process operation under state and input constraints. In *Proc. 40th IEEE Conf. Decision and Control*, pages 4228–4233, 2001.