Stochastic Model Predictive Control: State space methods

Mark Cannon[†]
Department of Engineering Science, University of Oxford,
mark.cannon@eng.ox.ac.uk

Contents

[†]These notes are based on research carried out at Oxford University jointly with Paul Couchman, Xingjian Wu, and Basil Kouvaritakis

1 Performance objective and closed-loop convergence

The performance objective of a Model Predictive Control algorithm determines the optimality, stability and convergence properties of the closed loop control law. In this section we consider how to generalize the quadratic cost typically employed in linear optimal control problems to account for stochastic model uncertainty. The emphasis is placed on a mean-variance formulation, which provides a trade-off between the objectives of minimizing the predicted performance of a nominal model and that of minimizing the variance of predicted future states and inputs. We derive expressions for the cost when evaluated over an infinite prediction horizon, and consider the closed loop properties of the optimal receding horizon controller. Under an assumption of generic soft constraints we discuss two convergence analyses, one based on bounds on l_2 gain, the other on a supermartingale property of the optimal value of predicted cost. Two types of convergence guarantees are obtained: asymptotic bounds on the time-average of the stage cost appearing in the performance objective, and repeated convergence of the system state to a region of state space.

1.1 Stochastic system models

Linear models for predicting time series data typically contain both auto-regressive (AR) and moving average (MA) components. Denoting the input (control variable) and system output sequences as $\{u_k,\ k = 0, 1, \ldots\}$ and $\{y_k,\ k = 0, 1, \ldots\}$, an nth order ARMA model has the form

$$y_k = \sum_{i=1}^{n} \alpha_{i,k}\, y_{k-i} + \sum_{i=i}^{n} \beta_{i,k}\, u_{k-i} + \gamma_k, \quad k = 0, 1, \ldots \tag{1.1}$$

with $\{\gamma_k\}$ denoting a zero-mean disturbance sequence. This form of model is typically identified using black-box methods, and is employed in diverse fields including process control and econometrics (Box and Jenkins, 1976; Brockwell and Davis, 2002). Although given here for the single-input single-output case, (1.1) is easily generalized to the case that u_k and y_k are vectors with n_u and n_y elements respectively.

The input-output response of the ARMA model in (1.1) could equivalently be generated using a state space realization of the system:

$$x_{k+1} = A_k x_k + B_k u_k + d_k, \quad y_k = C x_k \tag{1.2}$$

where x_k is the n-dimensional state vector. Matrices A, B, C depend linearly on α_i, β_i, and are non-unique in general (see e.g. Kailath, 1980). This section uses state space models because of their greater convenience for control design and estimation, and since since they are better suited to the dynamic programming recursions associated with cost evaluation and (in the next section) constraint handling over an infinite horizon.

If the disturbance γ and model coefficients α_i, β_i are random variables, then (1.1) defines very general form of linear stochastic system model. We assume that the model parameters α_i, β_i and disturbance γ are time-varying with stationary Gaussian distributions, and have means and covariance matrices that are known either from empirical identification algorithms or phenomenological modelling. Under these assumptions the parameters of the state space model (1.2) can be written as a linear expansion over a set of scalar random variables $\{q^{(1)}, \dots, q^{(m)}\}$:

$$\begin{bmatrix} A_k & B_k & d_k \end{bmatrix} = \begin{bmatrix} \bar{A} & \bar{B} & 0 \end{bmatrix} + \sum_{i=1}^{m} \begin{bmatrix} A^{(i)} & B^{(i)} & d^{(i)} \end{bmatrix} q_k^{(i)} \tag{1.3}$$

where $q_k = [q_k^{(1)} \ \cdots \ q_k^{(m)}]^T$ is a normally distributed random vector with zero mean and known covariance matrix S_q (denoted $q_k \sim \mathcal{N}(0, S_q)$). The following result shows that it is always possible to choose the basis $\{A^{(i)}, B^{(i)}, d^{(i)}\}$ over which the expansion of (1.3) is performed so that the elements of q_k are uncorrelated, and hence $S_q = I$ can be assumed without loss of generality.

Lemma 1.1. *There exist* $\{\hat{A}^{(i)}, \hat{B}^{(i)}, \hat{d}^{(i)}\}$ *satisfying*

$$\begin{bmatrix} A_k & B_k & d_k \end{bmatrix} = \begin{bmatrix} \bar{A} & \bar{B} & 0 \end{bmatrix} + \sum_{i=1}^{m} \begin{bmatrix} \hat{A}^{(i)} & \hat{B}^{(i)} & \hat{d}^{(i)} \end{bmatrix} \hat{q}_k^{(i)} \tag{1.4}$$

with $\hat{q}_k = [\hat{q}_k^{(1)} \ \cdots \ \hat{q}_k^{(m)}]^T \sim \mathcal{N}(0, I)$ *if* A, B, d *satisfy (1.3) and* $q_k \sim \mathcal{N}(0, S_q)$.

Proof. From (1.2) and (1.3) we have

$$x_{k+1} = \bar{A}x_k + \bar{B}u_k + \begin{bmatrix} A^{(1)}x_k + B^{(1)}u_k + d^{(1)} & \cdots & A^{(m)}x_k + B^{(m)}u_k + d^{(m)} \end{bmatrix} q_k.$$

Noting that S_q is necessarily symmetric and positive definite, let $S_q = VV^T$ and define $\hat{q}_k = V^{-1}q_k$. Then $\hat{q}_k \sim \mathcal{N}(0, I)$ and

$$x_{k+1} = \bar{A}x_k + \bar{B}u_k + \begin{bmatrix} A^{(1)}x_k + B^{(1)}u_k + d^{(1)} & \cdots & A^{(m)}x_k + B^{(m)}u_k + d^{(m)} \end{bmatrix} V\hat{q}_k$$

hence, with v_{ij} denoting the ijth element of V, we have

$$\begin{aligned} x_{k+1} &= \bar{A}x_k + \bar{B}u_k + \sum_{i=1}^{m}\sum_{j=1}^{m} \begin{bmatrix} A^{(i)} & B^{(i)} & d^{(i)} \end{bmatrix} \begin{bmatrix} x_k \\ u_k \\ 1 \end{bmatrix} v_{ij}\hat{q}_k^{(j)} \\ &= \bar{A}x_k + \bar{B}u_k + \sum_{j=1}^{m} \left(\sum_{i=1}^{m} \begin{bmatrix} A^{(i)} & B^{(i)} & d^{(i)} \end{bmatrix} v_{ij} \right) \begin{bmatrix} x_k \\ u_k \\ 1 \end{bmatrix} \hat{q}_k^{(j)} \end{aligned}$$

so that $\begin{bmatrix} \hat{A}^{(j)} & \hat{B}^{(j)} & \hat{d}^{(j)} \end{bmatrix} = \sum_{i=1}^{m} \begin{bmatrix} A^{(i)} & B^{(i)} & d^{(i)} \end{bmatrix} v_{ij}$ in (1.4). □

We make the further simplifying assumption that q_k and q_i are independent at all sampling instants $k \neq i$. Unlike the assumption that the expansion (1.3) is chosen so that the elements of q_k are uncorrelated, this property is not necessarily satisfied by the general linear stochastic model (1.2-1.3). However it can be ensured through a suitable augmentation of the state vector provided q_k and q_i are independent whenever $|k - i|$ is sufficiently large.

Given the value of the state x_k, the predictions of state and input trajectories based on x_k are denoted $\{x_{k+i|k}, u_{k+i|k},\ i = 0, 1, \dots\}$, with $x_{k|k} = x_k$. In practice x_k may not be directly measurable, in which case the measured state may be replaced by an observer estimate based on measurements of the output, y_k. We do not consider this output feedback problem explicitly here; instead we simply note that state estimation errors could be incorporated in the disturbance term of the prediction model (1.2).

Let r be a reference for the plant output y_k, and let $\bar{x}_{ss}$ be the expected value of the state satisfying the steady state conditions:

$$\bar{x}_{ss} = \bar{A}\bar{x}_{ss} + \bar{B}u_{ss}, \quad r = C\bar{x}_{ss}.$$

The problem of forcing y_k to track r is converted to one of regulating x_k to the origin if the following variable transformations are used:

$$\begin{aligned} x_k &\leftarrow x_k - \bar{x}_{ss} \\ u_k &\leftarrow u_k - u_{ss} \\ y_k &\leftarrow y_k - r \\ d_k &\leftarrow (B_k - \bar{B})u_{ss} + d_k = \sum_{i=1}^{m} (B^{(i)}u_{ss} + d^{(i)})q_k^{(i)} \end{aligned}$$

In the following we assume that these transformations have been applied to the state space model (1.2-1.3), and we therefore consider the aim of the controller to be regulation of x_k about the origin in a statistical sense.

1.2 Performance cost

The future predicted state trajectories of the model (1.2-1.3) are sequences of random variables, and perhaps the simplest way to account for this in the definition of a quadratic performance cost is to take expectations. Denoting the expectation operator conditional on information available at time k (namely the state x_k) as $\mathbb{E}_k$, consider the LQG cost (Astrom, 1970; Lee and Cooley, 1998):

$$J(x_k, \mathbf{u}_k) = \sum_{i=0}^{\infty} \mathbb{E}_k\big[\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2\big] \tag{1.5}$$

for $Q, R \succ 0$. Here $\mathbf{u}_k = \{u_{k|k}, u_{k+1|k}, \dots\}$ denotes a sequence of future control inputs predicted at time k and $\|x\|_Q^2 = x^T Q x$, $\|u\|_R^2 = u^T R u$.

To ensure that the stage cost $\mathbb{E}_k[\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2]$ converges to a finite limit as $i \to \infty$, we make the assumption that the pair (A, B) is mean-square stabilizable in the following sense.[1]

Assumption 1. *For any symmetric positive definite $S \in \mathbb{R}^{n_x \times n_x}$ there exists $K \in \mathbb{R}^{n_u \times n_x}$ and positive definite $P \in \mathbb{R}^{n_x \times n_x}$ satisfying*

$$P - \mathbb{E}\big[(A + BK)^T P (A + BK)\big] = S. \tag{1.6}$$

[1] See Kushner (1971) for a discussion of mean square stability

Given the model uncertainty (1.3), it is easy to show using Schur complements that condition (1.6) is satisfied for some K and P whenever the following LMI in variables Π and Γ is feasible:

$$\begin{bmatrix} \Pi & (\bar{A}\Pi + \bar{B}\Gamma)^T & (A^{(1)}\Pi + B^{(1)}\Gamma)^T & \cdots & (A^{(m)}\Pi + B^{(m)}\Gamma)^T & \Pi \\ \star & \Pi & 0 & \cdots & 0 & 0 \\ \star & \star & \Pi & \cdots & 0 & 0 \\ \vdots & \vdots & & \ddots & & \vdots \\ \star & \star & \star & \cdots & \Pi & 0 \\ \star & \star & \star & \cdots & 0 & S^{-1} \end{bmatrix} \succeq 0.$$

In the absence of additive disturbances (i.e. if $d_k = 0$ in (1.2) for all k), it can be shown (see e.g. Boyd et al., 1994) that Assumption 1 ensures that the optimal value of the quadratic cost $J(x_k, \mathbf{u}_k)$ is finite. As a result the optimal predicted trajectories for $x_{k+i|k}$ and $u_{k+i|k}$ would converge to zero as $i \to \infty$ with probability 1 (w.p.1) in this case (Kushner, 1971). However, for non-zero additive disturbance ($d_k \neq 0$), Assumption 1 implies that the stage cost in (1.5) converges to a non-zero value, and this in turn implies that $J(x_k, \mathbf{u}_k)$ is infinite. While this is not problematic for LQG control design (since the unconstrained optimal control can be obtained in closed form by solving a Riccati equation for the optimal linear feedback gain), it could cause difficulties for the computation of a predictive control law based on numerical optimization of $J(x_k, \mathbf{u}_k)$ subject to state and input constraints. To avoid this difficulty we redefine the MPC performance objective by subtracting the steady state value of the stage cost in (1.5). First however, it is necessary to determine this steady state value.

In order to define a predicted input sequence over an infinite horizon while requiring only a finite number of free variables to parameterize it, we make use of a dual mode prediction strategy (see e.g. Mayne et al., 2000). This specifies predicted inputs as a fixed linear feedback law at all times beyond an initial finite horizon of N steps:

$$u_{k+i|k} = K x_{k+i|k}, \quad i = N, N+1, \dots. \tag{1.7}$$

The steady state mean and variance of predicted state and input trajectories

are given by the following result.

Lemma 1.2. *Under the control law (1.7), $x_{k+i|k}$ and $u_{k+i|k}$ satisfy*

$$\lim_{i\to\infty} \mathbb{E}_k(x_{k+i|k}) = 0, \quad \lim_{i\to\infty} \mathbb{E}_k(u_{k+i|k}) = 0$$

and

$$\lim_{i\to\infty} \mathbb{E}_k\big[\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2\big] = \mathrm{tr}(\Theta L), \quad L = Q + K^T R K \tag{1.8a}$$

where $\Theta = \lim_{i\to\infty} \mathbb{E}_k(x_{k+i|k}x_{k+i|k}^T)$ is the positive definite solution of

$$\Theta - \mathbb{E}\big[(A+BK)\Theta(A+BK)^T\big] = \mathbb{E}(dd^T) \tag{1.8b}$$

if and only if K satisfies (1.6) for some $S, P \succ 0$.

Proof. The linearity of (1.2) implies that $x_{k+i|k} = \zeta_{k+i|k} + \xi_{k+i|k}$ for all $i \geq N$ where $\zeta_{k+i|k}$ and $\xi_{k+i|k}$ generated by the following pair of systems:

$$\zeta_{k+i+1|k} = \Phi_{k+i}\zeta_{k+i|k}, \qquad \zeta_{k+N|k} = x_{k+N|k} \tag{1.9a}$$

$$\xi_{k+i+1|k} = \Phi_{k+i}\xi_{k+i|k} + d_{k+i}, \qquad \xi_{k+N|k} = 0 \tag{1.9b}$$

in which $\Phi_k = A_k + B_k K$. Condition (1.6) is necessary and sufficient for mean square stability of Φ_k and (1.9a) therefore gives $\mathbb{E}_k(\zeta_{k+i|k}\zeta_{k+i|k}^T) \to 0$, and hence $\zeta_{k+i|k} \to 0$ w.p.1, as $i \to \infty$. From (1.9b) we have $\mathbb{E}_k(\xi_{k+i|k}) = 0$ for all $i \geq 0$, and it follows that $\mathbb{E}_k(x_{k+i|k}) \to 0$ and $\mathbb{E}_k(u_{k+i|k}) \to 0$ as $i \to \infty$. Noting also that $\xi_{k+i|k}$ is independent of Φ_{k+i}, (1.9b) gives

$$\begin{aligned}\mathbb{E}_k(\xi_{k+i+1|k}\xi_{k+i+1|k}^T) &= \mathbb{E}_k\left[(\Phi_{k+i}\xi_{k+i|k} + d_{k+i})(\Phi_{k+i}\xi_{k+i|k} + d_{k+i})^T\right] \\ &= \mathbb{E}_k(\Phi_{k+i}\xi_{k+i|k}\xi_{k+i|k}^T\Phi_{k+i}^T) + \mathbb{E}_k(d_{k+i}d_{k+i}^T)\end{aligned} \tag{1.10}$$

Let $\hat{\Theta}_{k+i|k} = \mathbb{E}_k(\xi_{k+i|k}\xi_{k+i|k}^T) - \Theta$, where Θ is the solution of (1.8b) then from (1.8b) and (1.10) we have

$$\begin{aligned}\hat{\Theta}_{k+i+1|k} &= \mathbb{E}_k(\Phi_{k+i}\xi_{k+i|k}\xi_{k+i|k}^T\Phi_{k+i}^T) - \mathbb{E}_k(\Phi_{k+i}\Theta\Phi_{k+i}^T) \\ &= \mathbb{E}_k(\Phi_{k+i}\hat{\Theta}_{k+i|k}\Phi_{k+i}^T).\end{aligned}$$

The mean square stability of Φ_k therefore implies that $\hat{\Theta}_{k+i|k} \to 0$, so that $\mathbb{E}_k(\xi_{k+i|k}\xi_{k+i|k}^T) \to \Theta$ as $i \to \infty$. But $\mathbb{E}_k(x_{k+i|k}x_{k+i|k}^T) \to \mathbb{E}_k(\xi_{k+i|k}\xi_{k+i|k}^T)$ since $\zeta_{k+i|k} \to 0$ as $i \to \infty$, and thus $\lim_{i\to\infty} \mathbb{E}_k(x_{k+i|k}x_{k+i|k}^T) = \Theta$. □

On the basis of Lemma 1.2, we redefine the LQG cost (1.5) as

$$J(x_k, \mathbf{u}_k) = \sum_{i=0}^{\infty}\big\{\mathbb{E}_k\big[\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2\big] - \mathrm{tr}(\Theta L)\big\}. \tag{1.11}$$

This performance objective has several desirable properties:

- computational convenience — in Section 1.3 we show that it can be expressed as a quadratic function of the degrees of freedom in predictions
- a closed-form solution for the optimal $\mathbf{u}_k$ in the absence of constraints (which is identical to the optimal feedback law for the original LQG cost (1.5))
- the optimal value can be used to define a stochastic Lyapunov function when used in a receding horizon control law (as discussed in Section 1.5).

However the LQG cost (1.11) optimizes just one aspect of the distribution of predicted trajectories, namely the second moment. Moreover it is generally desirable to use a cost that can trade off conflicting requirements for nominal performance (computed using a nominal model, such as the model that is obtained by setting all uncertain parameters to their expected values) against the worst case performance that corresponds a specified confidence level on model parameters. This is the motivation for the mean-variance cost proposed in Couchman et al. (2006b) and Cannon et al. (2007), which is of the form

$$\begin{aligned}V(x_k, \mathbf{u}_k) = &\sum_{i=0}^{\infty}\big[\|\bar{x}_{k+i|k}\|_Q^2 + \|\bar{u}_{k+i|k}\|_R^2\big] \\ &+ \kappa^2\sum_{i=0}^{\infty}\big\{\mathbb{E}_k\big[\|x_{k+i|k} - \bar{x}_{k+i|k}\|_Q^2 + \|u_{k+i|k} - \bar{u}_{k+i|k}\|_R^2\big] - \mathrm{tr}(\Theta L)\big\}.\end{aligned} \tag{1.12}$$

Here $\bar{x}_{k+i|k} = \mathbb{E}_k(x_{k+i|k})$ and $\bar{u}_{k+i|k} = \mathbb{E}_k(u_{k+i|k})$, so that

$$\bar{x}_{k+i+1|k} = \bar{A}\bar{x}_{k+i|k} + \bar{u}_{k+i|k},$$

and κ is a design variable that controls the relative weighting of mean and variance terms in the cost.

If $A_k = \bar{A}$ in (1.3), so that the uncertainty in the model parameters is restricted to the input map B_k and additive disturbance d_k, then the state predictions

$x_{k+i|k}$ are normally distributed for all i. If, in addition, $Q = CC^T$ and $R = 0$, so that $\|x\|_Q^2 = y^2$, then the parameter κ can be interpreted in terms of probabilistic bounds on the predicted output trajectory $\{y_{k+i|k},\ i = 1, 2, \ldots\}$ (Cannon et al., 2007). To see this, let $\lambda_{k+i|k}$ and $\upsilon_{k+i|k}$ be lower and upper bounds on $y_{k+i|k}$ that hold with a given probability p:

$$\begin{aligned} \Pr(y_{k+i|k} \geq \lambda_{k+i|k}) \geq p \\ \Pr(y_{k+i|k} \leq \upsilon_{k+i|k}) \geq p \end{aligned} \tag{1.13}$$

then, for the case that $y_{k+i|k}$ is normally distributed it is straightforward to show that

$$\bar{y}_{k+i|k}^2 + \kappa^2 \mathbb{E}_k (y_{k+i|k} - \bar{y}_{k+i|k})^2 = \tfrac{1}{2}\lambda_{k+j|k}^2 + \tfrac{1}{2}\upsilon_{k+j|k}^2$$

provided κ satisfies $\mathfrak{N}(\kappa) = p$, where $\mathfrak{N}$ is the normal distribution function: $\Pr(z \leq Z) = \mathfrak{N}(Z)$ for $z \sim \mathcal{N}(0, 1)$.

For the general case in which A_k contains uncertain parameters, the state predictions $x_{k+i|k}$ are not normally distributed for $i \geq 2$, and κ does not therefore have an interpretation in terms of probabilistic bounds on predictions. However the cost (1.12) does provide a means of balancing nominal and minimum-variance performance objectives, and this form of objective has been proposed for applications such as optimal portfolio selection (Zhu et al., 2004) and sustainable development problems (Couchman et al., 2006a,b). We allow κ to take any (fixed) value in the range $0 \leq \kappa < \infty$ in (1.12), and the special cases of:

- the nominal cost if $\kappa = 0$
- the LQG cost (1.5) if $\kappa = 1$
- minimum variance cost in the limit as $\kappa \to \infty$

are therefore included in this framework.

1.3 Cost evaluation

This section considers how to express the cost (1.12) as a function of the degrees of freedom in state and input trajectories in order to obtain the objective function in a form that can be optimized online.

We define the predicted trajectory for $u_{k+i|k}$ using the closed loop paradigm of Kouvaritakis et al. (2000), which embeds the feedback law of (1.7) into input predictions at all future times:

$$u_{k+i|k} = K x_{k+i|k} + f_{i|k}, \quad i = 0, 1, \ldots \tag{1.14a}$$

$$f_{i|k} = 0, \quad i = N, N+1, \ldots \tag{1.14b}$$

Here K is assumed to satisfy the mean square stability condition (1.6) in Assumption 1, and ideally K should be optimal in the absence of input and state constraints. Though infinite, the input sequence defined by (1.14) contains a finite number of free variables $f_{i|k}$, $i = 0, \ldots, N-1$, making numerical optimization of a predicted cost practicable.

For the class of predicted inputs defined by (1.14), the dynamics (1.2) give the predicted state trajectories as

$$x_{k+i+1|k} = \Phi_{k+i} x_{k+i|k} + B_{k+i} f_{i|k} + d_{k+i}, \quad \Phi_k = A_k + B_k K. \tag{1.15}$$

To express the performance cost as a function of $\mathbf{f}_k = \begin{bmatrix} f_{0|k}^T & \cdots & f_{N-1|k}^T \end{bmatrix}^T$ we use an autonomous formulation of the prediction dynamics, which generates the predictions of (1.14-1.15) via:

$$z_{k+i+1|k} = \Psi_{k+i} z_{k+i|k}, \quad z_{k|k} = \begin{bmatrix} x_k \\ \mathbf{f}_k \\ 1 \end{bmatrix}, \quad \Psi_k = \begin{bmatrix} \Phi_k & B_k E & d_k \\ 0 & M & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{1.16a}$$

$$M = \begin{bmatrix} 0 & I_{n_u} & 0 & \ldots & 0 \\ 0 & 0 & I_{n_u} & & 0 \\ \vdots & \vdots & & & \vdots \\ 0 & 0 & 0 & \ldots & 0 \end{bmatrix}, \quad E = \begin{bmatrix} I_{n_u} & 0 & \ldots & 0 \end{bmatrix} \tag{1.16b}$$

with $x_{k+i|k} = \begin{bmatrix} I_{n_x} & 0 & 0 \end{bmatrix} z_{k+i|k}$, $u_{k+i|k} = \begin{bmatrix} K & E & 0 \end{bmatrix} z_{k+i|k}$.

This formulation is autonomous in respect of the control input over the prediction horizon (although it retains the exogenous multiplicative and additive disturbances present in (1.2-1.3)). It can be interpreted as dynamic feedback

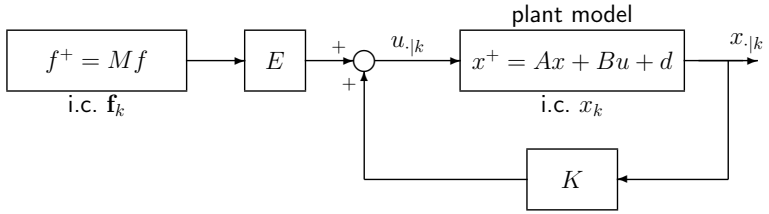

Figure 1: Block diagram representation of the autonomous prediction system (1.16a). The optimization variables are the initial controller state $\mathbf{f}_k$

law applied to (1.2), with the degrees of freedom in predictions contained in the controller state at the beginning of the prediction horizon (see Fig. 1). The advantage of this formulation is that it enables the two components of the cost (1.12) to be computed by solving a pair of Lyapunov matrix equations, as described in the following theorem. By a slight abuse of notation we denote V evaluated along trajectories of the prediction system (1.16a) as $V(x_k, \mathbf{f}_k)$.

Theorem 1.3. *Along trajectories of* (1.16a) *the cost* (1.12) *is given by*

$$V(x_k, \mathbf{f}_k) = z_{k|k}^T P z_{k|k}, \quad P = \begin{bmatrix} P_x & P_{xf} & P_{x1} \\ P_{fx} & P_f & P_{f1} \\ P_{1x} & P_{1f} & P_1 \end{bmatrix} \tag{1.17a}$$

$$P = (1-\kappa^2)X + \kappa^2 Y \tag{1.17b}$$

where X, Y are partitioned into blocks that are conformal to the partition of P in (1.17a) and defined by

$$\begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} - \begin{bmatrix} \bar\Phi & \bar B E \\ 0 & M \end{bmatrix}^T \begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} \begin{bmatrix} \bar\Phi & \bar B E \\ 0 & M \end{bmatrix} = \begin{bmatrix} L & K^T R E \\ E^T R K & E^T R E \end{bmatrix} \tag{1.18a}$$

$$\begin{bmatrix} X_{1x} & X_{1f} \end{bmatrix} = \begin{bmatrix} 0 & 0 \end{bmatrix}, \quad X_1 = 0 \tag{1.18b}$$

and

$$\begin{bmatrix} Y_x & Y_{xf} \\ Y_{fx} & Y_f \end{bmatrix} - \mathbb{E}\left(\begin{bmatrix} \Phi & BE \\ 0 & M \end{bmatrix}^T \begin{bmatrix} Y_x & Y_{xf} \\ Y_{fx} & Y_f \end{bmatrix} \begin{bmatrix} \Phi & BE \\ 0 & M \end{bmatrix} \right) = \begin{bmatrix} L & K^T R E \\ E^T R K & E^T R E \end{bmatrix} \tag{1.19a}$$

$$\begin{bmatrix} Y_{1x} & Y_{1f} \end{bmatrix} \left(I - \begin{bmatrix} \bar\Phi & \bar B E \\ 0 & M \end{bmatrix} \right) = \mathbb{E}\Big(d^T Y_x \begin{bmatrix} \Phi & BE \end{bmatrix} \Big) \tag{1.19b}$$

$$Y_1 = -\mathrm{tr}(\Theta Y_x) \tag{1.19c}$$

with $\Phi = A + BK$ and $\bar\Phi = \bar A + \bar B K$.

Proof. Using the identity $\mathbb{E}(\|x - \bar x\|_Q^2 = \mathbb{E}(\|x\|_Q^2) - \|\bar x\|_Q^2$, $V(x_k, \mathbf{f}_k)$ can be rewritten as:

$$\begin{aligned} V(x_k, \mathbf{f}_k) = (1-\kappa^2) \sum_{i=0}^{\infty} \big[\|\bar x_{k+i|k}\|_Q^2 + \|\bar u_{k+i|k}\|_R^2 \big] \\ + \kappa^2 \sum_{i=0}^{\infty} \big\{ \mathbb{E}_k \big[\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2 \big] - \mathrm{tr}(\Theta L) \big\}. \end{aligned} \tag{1.20}$$

Consider the first sum on the RHS of (1.20). Taking expectations of (1.16a) and using (1.18) gives

$$\bar z_{k+i|k}^T X \bar z_{k+i|k} - \bar z_{k+i+1|k}^T X \bar z_{k+i+1|k} = \|\bar x_{k+i|k}\|_Q^2 + \|\bar u_{k+i|k}^2\|_R^2.$$

Summing both sides of this equation over all $i \geq 0$ and noting that Lemma 1.2 implies that $\lim_{i\to\infty} \bar z_{k+i|k} = 0$, it follows that

$$z_{k|k}^T X z_{k|k} = \sum_{i=0}^{\infty} \big[\|\bar x_{k+i|k}\|_Q^2 + \|\bar u_{k+i|k}\|_R^2 \big]. \tag{1.21}$$

Next consider the second sum in (1.20). From (1.16a) we have

$$\begin{aligned} & z_{k+i|k}^T Y z_{k+i|k} - \mathbb{E}_k(z_{k+i+1|k}^T Y z_{k+i+1|k}) = \\ & \begin{bmatrix} x_{k+i|k} \\ \mathbf{f}_{k+i|k} \end{bmatrix}^T \left\{ \begin{bmatrix} Y_x & Y_{xf} \\ Y_{fx} & Y_f \end{bmatrix} - \mathbb{E}\left(\begin{bmatrix} \Phi & BE \\ 0 & M \end{bmatrix}^T \begin{bmatrix} Y_x & Y_{xf} \\ Y_{fx} & Y_f \end{bmatrix} \begin{bmatrix} \Phi & BE \\ 0 & M \end{bmatrix} \right) \right\} \begin{bmatrix} x_{k+i|k} \\ \mathbf{f}_{k+i|k} \end{bmatrix} \\ & + 2 \left\{ \begin{bmatrix} Y_{1x} & Y_{1f} \end{bmatrix} \left(I - \begin{bmatrix} \bar\Phi & \bar B E \\ 0 & M \end{bmatrix} \right) - \mathbb{E}\Big(d^T Y_x \begin{bmatrix} \Phi & BE \end{bmatrix} \Big) \right\} \begin{bmatrix} x_{k+i|k} \\ \mathbf{f}_{k+i|k} \end{bmatrix} \\ & - \mathbb{E}(d^T Y_x d) \end{aligned}$$

and using (1.19) therefore gives

$$z_{k+i|k}^T Y z_{k+i|k} - \mathbb{E}_k(z_{k+i+1|k}^T Y z_{k+i+1|k}) = \mathbb{E}_k(\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2) - \mathbb{E}(d^T Y_x d). \tag{1.22}$$

Furthermore, from equation (1.8b) defining Θ we have

$$\operatorname{tr}\big(\Theta Y_x - \mathbb{E}(\Phi\Theta\Phi^T)Y_x\big) = \operatorname{tr}\big(\Theta[Y_x - \mathbb{E}(\Phi^T Y_x \Phi^T)]\big) = \operatorname{tr}\big(\mathbb{E}(dd^T)Y_x\big)$$

so that (1.19a) implies

$$\mathbb{E}(d^T Y_x d) = \operatorname{tr}(\Theta L) \tag{1.23}$$

which, when substituted into (1.22) gives

$$z_{k+i|k}^T Y z_{k+i|k} - \mathbb{E}_k(z_{k+i+1|k}^T Y z_{k+i+1|k}) = \mathbb{E}_k(\|x_{k+i|k}\|_Q^2 + \|u_{k+i|k}\|_R^2) - \operatorname{tr}(\Theta L).$$

Summing both sides over all $i \geq 0$ and noting that $\lim_{i\to\infty} \mathbb{E}_k(z_{i|k}^T Y z_{i|k}) = 0$ (since $\mathbf{f}_{k+i|k} = 0$ for $i \geq N$, while from Lemma 1.2 $\lim_{i\to\infty} \mathbb{E}_k(x_{i|k}) = 0$ and $\lim_{i\to\infty} \mathbb{E}_k(x_{i|k}x_{i|k}^T) = \Theta$, so that $\mathbb{E}_k(z_{i|k}^T Y z_{i|k}) \to \mathbb{E}_k(x_{i|k}^T Y_x x_{i|k}) - \operatorname{tr}(\Theta Y_x) \to 0$ as $i \to \infty$), we therefore have

$$z_{k|k}^T Y z_{k|k} = \sum_{i=0}^{\infty} \big\{\mathbb{E}_k\big[\|\bar{x}_{k+i|k}\|_Q^2 + \|\bar{u}_{k+i|k}\|_R^2\big] - \operatorname{tr}(\Theta L)\big\}. \tag{1.24}$$

From (1.21) and (1.24) it follows that the cost (1.12) is given by (1.17). □

1.4 Unconstrained optimal control

In this section we discuss how to compute an optimal value for the gain K in predictions (1.14-1.15). We show that, in the absence of input and state constraints, or whenever such constraints are inactive, the minimizing control for (1.12) is an affine state feedback law with a linear gain matrix defined by a pair of coupled algebraic Riccati equations (CAREs). In the interests of optimality, and to ensure the closed-loop convergence properties given in Section 1.5, we define K as this unconstrained optimal feedback gain. An iterative method for solving for K is described, which is similar to the Lyapunov iterations proposed for CAREs associated with related continuous-time control problems (Gajic and Borno, 1995; Freiling et al., 1999). We give necessary

and sufficient conditions for convergence of the iteration to the unconstrained optimal feedback gain.

To determine the unconstrained optimal feedback policy we define the following dynamic programming problem.

$$\begin{aligned}\mathbb{V}_k(x_k, u_k, i) = \kappa^2\big(\|x_k\|_Q^2 + \|u_k\|_R^2\big) + (1-\kappa^2)\big[\|\mathbb{E}_i(x_k)\|_Q^2 + \|\mathbb{E}_i(u_k)\|_R^2\big] \\ + \mathbb{E}_i\big[\mathbb{V}_{k+1}(x_{k+1}, u_{k+1}^*, i)\big]\end{aligned} \tag{1.25a}$$

$$u_k^*(x_k) = \arg\min_{u_k} \mathbb{V}_k(x_k, u_k, k) \tag{1.25b}$$

This is connected to the performance index in (1.12) by the fact that $\mathbb{V}(x_k, u_k, k)$ is equivalent to $V(x_k, \mathbf{u}_k)$ if $\mathbf{u}_k$ is defined by the infinite sequence $\{u_k^*, u_{k+1}^* \ldots\}$. Unlike the dynamic programming problem associated with the LQG cost (1.5), Bellman's principle of optimality (Bellman, 1952) does not hold for this problem. This is because the expected value of the optimal cost-to-go at the kth stage is not equal to the expectation of the $k+1$th-stage optimal cost, i.e.

$$\mathbb{E}_k\big[\mathbb{V}_{k+1}(x_{k+1}, u_{k+1}^*, k)\big] \neq \mathbb{E}_k\big[\mathbb{V}_{k+1}(x_{k+1}, u_{k+1}^*, k+1)\big].$$

However it is possible to use dynamic programming recursions to determine the optimal control law, as demonstrated next.

Lemma 1.4. *The optimal value function and optimal control law in (1.25) are given by*

$$\begin{aligned}\mathbb{V}_k(x_k, u_k^*, k) = \kappa^2(x_k^T Y_x^{(k)} x_k + 2Y_{1x}^{(k)} x_k + Y_1^{(k)}) \\ + (1-\kappa^2)(x_k^T X_x^{(k)} x_k + 2X_{1x}^{(k)} x_k + X_1^{(k)})\end{aligned} \tag{1.26a}$$

$$u_k^*(x_k) = K^{(k)} x_k + l^{(k)} \tag{1.26b}$$

where

$$K^{(k)} = -(D^{(k)})^{-1}\big[\kappa^2\mathbb{E}(B^T Y_x^{(k+1)} A) + (1-\kappa^2)\bar{B}^T X_x^{(k+1)} \bar{A}\big] \tag{1.27a}$$

$$l^{(k)} = -(D^{(k)})^{-1}\big[\kappa^2\mathbb{E}(B^T Y_x^{(k+1)} d) + \kappa^2 \bar{B}^T Y_{x1}^{(k+1)} + (1-\kappa^2)\bar{B}^T X_{x1}^{(k+1)}\big] \tag{1.27b}$$

with $D^{(k)} = R + \kappa^2\mathbb{E}(B^T Y_x^{(k+1)} B) + (1-\kappa^2)\bar{B}^T X_x^{(k+1)} \bar{B}$. *Here* $Y_x^{(k)}, Y_{x1}^{(k)}, Y_1^{(k)}$

and $X_x^{(k)}, X_{x1}^{(k)}, X_1^{(k)}$ are defined by recursion relations, and specifically:

$$Y_x^{(k)} = \mathbb{E}\big[(A+BK^{(k)})^T Y_x^{(k+1)}(A+BK^{(k)})\big] + Q + K^{(k)T}RK^{(k)} \quad (1.28a)$$
$$X_x^{(k)} = (\bar{A}+\bar{B}K^{(k)})^T X_x^{(k+1)}(\bar{A}+\bar{B}K^{(k)}) + Q + K^{(k)T}RK^{(k)}. \quad (1.28b)$$

Proof. The quadratic form in (1.26a) can be shown by induction on k, and the affine control law (1.26b) is a direct consequence of this. □

Before considering the convergence of the dynamic programming recursions given in (1.28), we first discuss the link between the optimal control in (1.26b) and the unconstrained solution to the problem of minimizing $V(x_k, \mathbf{f}_k)$.

Lemma 1.5. *The minimizer $\mathbf{f}^*(x_k) = \arg\min_{\mathbf{f}} V(x_k, \mathbf{f})$ is unique and satisfies*

$$\mathbf{f}^*(x_k) = -P_x^{-1}P_{fx}x_k - P_x^{-1}P_{f1}. \quad (1.29)$$

Proof. For $\Gamma = \begin{bmatrix} K & E \end{bmatrix}$, (1.18a) implies

$$\begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} - \bar{\Psi}_{xf}^T \begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} \bar{\Psi}_{xf} \succeq \Gamma\Gamma^T, \quad \bar{\Psi}_{xf} = \begin{bmatrix} \bar{\Phi} & \bar{B}E \\ 0 & M \end{bmatrix}$$

where, by construction, $\bar{\Psi}_{xf}$ is stable and $(\bar{\Psi}_{xf}, \Gamma)$ is observable. From a well-known property of Lyapunov matrix equations, it follows that

$$\begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} \succ 0. \quad (1.30)$$

Similarly, subtracting (1.18a) from (1.19a) gives

$$\begin{bmatrix} Y_x - X_x & Y_{xf} - X_{xf} \\ Y_{fx} - X_{fx} & Y_f - X_f \end{bmatrix} - \bar{\Psi}_{xf}^T \begin{bmatrix} Y_x - X_x & Y_{xf} - X_{xf} \\ Y_{fx} - X_{fx} & Y_f - X_f \end{bmatrix} \bar{\Psi}_{xf} \succeq 0$$

which implies that

$$\begin{bmatrix} Y_x & Y_{xf} \\ Y_{fx} & Y_f \end{bmatrix} - \begin{bmatrix} X_x & X_{xf} \\ X_{fx} & X_f \end{bmatrix} \succeq 0. \quad (1.31)$$

From (1.30), (1.31) and $P = X + \kappa^2(Y - X)$, it follows that $P_f \succ 0$ for all $\kappa > 0$ and all $N \geq 1$. Hence $\mathbf{f}^*(x_k)$ is defined uniquely by (1.29). □

The optimal $\mathbf{f}_k^*$ given in (1.29) is the solution of an open-loop optimal control problem, since the elements $f_{i|k}^*$ of $\mathbf{f}_k^*$ are in general functions of x_k rather than functions of $x_{k+i|k}$. However an exception to this is the special case in which $\mathbf{f}_k^*$ is independent of x_k. The following theorem uses this observation to show that, if $\mathbf{f}_k^*$ is independent of x_k, then the predicted input trajectory corresponding to the optimal open-loop solution $\mathbf{f}_k^*$ for $N = \infty$ is identical to the optimal feedback law in (1.25) for the infinite horizon case, so that $u_{k+i|k} = u_{k+i}^*(x_{k+i|k})$.

Theorem 1.6. *The cost (1.12) is minimized over all input sequences $\mathbf{u}_k$ by the affine state feedback control law $u_{k+i|k} = Kx_{k+i|k} + l$ with*

$$K = -D^{-1}\big[(1-\kappa^2)\bar{B}^T X_x \bar{A} + \kappa^2 \mathbb{E}(B^T Y_x A)\big] \quad (1.32a)$$
$$l = -D^{-1}\kappa^2\big[\mathbb{E}(B^T Y_x d) + \bar{B}^T Y_{x1}\big] \quad (1.32b)$$

and $D = R + (1-\kappa^2)\bar{B}^T X_x \bar{B} + \kappa^2 \mathbb{E}(B^T Y_x B)$.

Proof. From (1.29), $\mathbf{f}^*$ is independent of x_k iff $P_{fx} = 0$. But

$$P_{fx} - M^T P_{fv}\bar{\Phi} = E^T\big[RK + (1-\kappa^2)\bar{B}^T X_x \bar{\Phi} + \kappa^2 \mathbb{E}(B^T Y_x \Phi)\big]$$

from (1.19a), and hence $P_{fx} = 0$ for all N if and only if K satisfies (1.32a). Under this condition (1.19b) implies that each element of $\mathbf{f}^*$ is equal to l defined in (1.32b). □

Comparing (1.32) with (1.27) and comparing the definitions of X_x, Y_x in (1.18-1.19) with the recursion (1.28), it is clear that the optimal feedback gain in Theorem 1.6 corresponds to a fixed point of the iteration (1.28). This can be stated as follows.

Corollary 1.7. *The optimal feedback gain in (1.32a) is defined by the solution of the coupled algebraic Riccati equations (CARE):*

$$X_x = (\bar{A}+\bar{B}K)^T X_x (\bar{A}+\bar{B}K) + Q + K^T R K \quad (1.33a)$$
$$Y_x = \mathbb{E}\big[(A+BK)^T Y_x (A+BK)\big] + Q + K^T R K \quad (1.33b)$$

where $K = -D^{-1}[(1-\kappa^2)\bar{B}^T X_x \bar{A} + \kappa^2 \mathbb{E}(B^T Y_x A)]$.

The problem of solving (1.33) for (X_x, Y_x, K) is non-convex and, unlike the conventional algebraic Riccati equation, there is no transformation that leads to a LMI formulation. We therefore propose to compute (X_x, Y_x, K) using the dynamic programming iteration given in Lemma 1.4.

Algorithm 1. Set $X_x^{(0)} = Y_x^{(0)} = Q$. For $k = -1, -2 \ldots$, compute $K^{(k)}$ using (1.27a) and $Y_x^{(k)}, X_x^{(k)}$ using (1.28a,b). Stop when $\|K^{(k)} - K^{(k+1)}\| < \epsilon$ for some suitably small tolerance ϵ, and set $K = K^{(k)}$.

It is not immediately obvious that the iteration employed in Algorithm 1 converges since, as mentioned previously, the principle of optimality does not apply to the dynamic programming problem in (1.25). However it is possible to show that Algorithm 1 converges in a finite number of iterations using the following theorem.

Theorem 1.8. *The sequence* $\{K^{(k)}, X_x^{(k)}, Y_x^{(k)},\ k = 0, -1, -2 \ldots\}$ *generated by the iteration (1.27a-1.28a,b) converges as* $k \to -\infty$ *to a fixed point satisfying (1.33) if and only if* (A, B) *satisfies the stabilizability condition of Assumption 1.*

A sketch of the proof is as follows.

1. It can be shown that $X_x^{(k)}$ and $Y_x^{(k)}$ are bounded for all k if and only if Assumption 1 holds. This is a consequence of $K^{(k)}$ in (1.27a) satisfying

$$K^{(k)} = \arg\min_{K^{(k)}} \mathrm{tr}\big[(1-\kappa^2)X_x^{(k)} + \kappa^2 Y_x^{(k)}\big] \quad \text{subject to (1.28a,b)}$$

   and the fact that, for any $\kappa \geq 0$, it is always possible to choose $K^{(k)}$ so that $\mathrm{tr}[(1-\kappa^2)X_x^{(k)} + \kappa^2 Y_x^{(k)}] < \mathrm{tr}[(1-\kappa^2)X_x^{(k+1)} + \kappa^2 Y_x^{(k+1)}]$ whenever $\mathrm{tr}(X_x^{(k+1)})$ or $\mathrm{tr}(Y_x^{(k+1)})$ are sufficiently large.

2. From the Bolzano-Weierstrass theorem, it then follows that the sequence $\{K^{(k)}, X_x^{(k)}, Y_x^{(k)},\ k = 0, -1, -2 \ldots\}$ has at least one accumulation point Kreyszig (2006). Given the autonomous nature of the iteration (1.27a-1.28a,b), the sequence $K^{(k)}$ therefore must converge as $k \to -\infty$ either to a fixed point satisfying (1.33) or to a limit cycle $\{K^{(k)},\ k = 1, \ldots, n\}$, with $K^{(1)} = K^{(n)}$. The latter case can be ruled out since it would imply that the optimal input sequence for (1.12) is one of n periodic feedback laws, and would thus contradict the uniqueness result of Lemma 1.5.

1.5 Receding horizon control, stability and convergence

The unconstrained optimal control law described in Section 1.4 will not meet constraints in general, and it becomes necessary instead to approach the constrained optimization of performance in a receding horizon manner, by minimizing the predicted cost online over the vector of degrees of freedom $\mathbf{f}_k$, subject to constraints. This section defines a general MPC law for systems with soft constraints, demonstrates that a stochastic form of stability holds for the resulting closed loop system, and derives bounds on the convergence of the time-average of the state covariance.

Given that no bounds have been assumed for the stochastic parameter q_k in (1.3), it is unrealistic to expect that state constraints will be satisfied for all realizations of uncertainty along predicted trajectories. Instead, we consider in this section the case of soft input and state constraints which are stated in terms of the nominal model dynamics:

$$F\bar{x}_{k+i|k} + G\bar{u}_{k+i|k} \leq h, \quad i = 0, 1, \ldots \tag{1.34}$$

Let $\mathcal{F}$ denote the feasible region in which x_k must lie so that (1.34) can be met over the prediction horizon for some $\mathbf{f}_k$. Given the linearity of the model (1.2) and constraints (1.34), $\mathcal{F}$ is a polytopic set, which can computed using the approach of Gilbert and Tan (1991), or otherwise can be inner-approximated by an ellipsoidal set which is invariant with respect to the nominal dynamics (e.g. the projection onto the x-subspace of an augmented invariant set in $(x, \mathbf{f})$-space proposed in Kouvaritakis et al. (2000)). We define the following receding horizon control law

Algorithm 2. At times $k = 0, 1, \ldots$:
1. If $x_k \in \mathcal{F}$, minimize the cost (1.17) by computing

$$\mathbf{f}^*(x_k) = \arg\min_{\mathbf{f}} V(x_k, \mathbf{f}) \quad \text{subject to (1.34)} \tag{1.35}$$

2. If $x_k \notin \mathcal{F}$, minimize the maximum constraint violation via

$$\mathbf{f}^*(x_k) = \arg\min_{\mathbf{f}} \max_i F\bar{x}(k+i|k) + G\bar{u}(k+i|k) - h$$
$$\text{subject to } V(x_k, \mathbf{f}) \le V(x_k, M\mathbf{f}^*(x_{k-1})) \quad (1.36)$$

3. Implement $u_k = Kx_k + E\mathbf{f}^*(x_k)$.

Note that (1.35) and (1.36) are convex, and that (1.35) is a linearly constrained quadratic program, whereas (1.36) is a quadratically constrained linear program. For both of these problems efficient solvers are available. Furthermore, due to the definition of V, the constraint in (1.36) is necessarily feasible. The purpose of this constraint is to enforce the property that the expected value of $V(x_k, \mathbf{f}^*(x_k))$ decreases whenever the state is sufficiently large.

For the general case of $\kappa \ge 0$, stability is analyzed using input-state stability arguments, whereas a supermartingale analysis is used in Section 1.6 for the special case of $\kappa \ge 1$. The input-state stability approach requires knowledge of the structure of some of the blocks of P in (1.17), and, under the assumption that K is the unconstrained optimal feedback gain, this is provided in the result below.

Lemma 1.9. *If K satisfies (1.32a), then the blocks P_{fx}, P_f, P_{f1} of P are given by*

$$P_{xf} = 0, \quad P_f = \operatorname{diag}\{D \ \dots \ D\}, \quad P_{f1} = [g^T \ \dots \ g^T]^T$$

where D is as defined in Theorem 1.6 and $g = \kappa^2\big[\mathbb{E}(B^T Y_x d) + \bar{B}^T Y_{x1}\big]$.

Proof. $P_{xf} = 0$ follows directly from Theorem 1.6. The block structure of P_f and P_{f1} is due to the structure of E and M in the prediction dynamics (1.16a); for brevity we omit the derivation. □

In the sequel it is assumed that K is the unconstrained optimal feedback gain, and the cost $V(x_k, \mathbf{f}_k)$ can therefore be expressed:

$$V(x_k, \mathbf{f}_k) = V_f(\mathbf{f}_k) + V_x(x_k) \quad (1.37a)$$
$$V_f(\mathbf{f}) = \mathbf{f}^T P_f \mathbf{f} + 2\mathbf{f}^T P_{f1} \quad (1.37b)$$
$$V_x(x) = x^T P_x x + 2x^T P_{x1} + P_1 \quad (1.37c)$$

The following stability analysis derives bounds on $f^*_{0|k}$ and uses these to bound a quadratic function of the state. To simplify analysis we consider only $V_f(\mathbf{f}_k)$, with the understanding that, for given x_k, minimizing V is equivalent to minimizing V_f.

Lemma 1.10. *Let $V^*_{f,k} = V_f\big(\mathbf{f}^*(x_k)\big)$. Then $\mathbf{f}^*$ in (1.35) satisfies*

$$\sum_{k=0}^{n-1} \mathbb{E}_0\Big[\|f^*_{0|k}\|^2_D + 2g^T f^*_{0|k}\Big] \le V^*_{f,0} - \mathbb{E}_0(V^*_{f,n-1}). \quad (1.38)$$

Proof. From the objective in (1.35) and the constraint in (1.36), we obtain $V(x_k, \mathbf{f}^*_k) \le V(x_k, M\mathbf{f}^*_{k-1})$ which implies that $V_f(\mathbf{f}^*_k) \le V_f(M\mathbf{f}^*_{k-1})$. But Lemma 1.9 and (1.37) imply $V_f(M\mathbf{f}^*_{k-1}) = V_f(\mathbf{f}^*_{k-1}) - \|f^*_{0|k-1}\|^2_D - 2g^T f^*_{0|k}$ and therefore

$$\mathbb{E}_{k-1}\big[V_f(\mathbf{f}^*_k)\big] \le V_f(\mathbf{f}^*_{k-1}) - \|f^*_{0|k-1}\|^2_D - 2g^T f^*_{0|k}. \quad (1.39)$$

The inequality (1.38) is obtained by applying the operator $\mathbb{E}_0$ to both sides of (1.39), using the property that $\mathbb{E}_0\big[\mathbb{E}_i[\cdot]\big] = \mathbb{E}_0[\cdot]$ for all $i > 0$, and summing over $k = 1, \dots, n$. □

We next show that the dynamics of the closed loop system mapping $f^*_{0|k}$ to x_k have finite l_2-gain.

Theorem 1.11. *For any given $\hat{g} \in \mathbb{R}^{n_u}$, there exist $\beta > 0$ and $\hat{Y}_x \succ 0$ such that*

$$\sum_{k=0}^{n-1}\big(\mathbb{E}_0\big[x_k^T L x_k\big] - \operatorname{tr}(\Theta L)\big) \le \sum_{k=0}^{n-1} \mathbb{E}_0\big[\beta\|f^*_{0|k}\|^2 + 2\hat{g}^T f^*_{0|k}\big] + x_0^T \hat{Y}_x x_0 - \mathbb{E}_0\big[x_n^T \hat{Y}_x x_n\big]. \quad (1.40)$$

Proof. Given that $\Phi = A + BK$ is mean square stable, let $\hat{Y}_x = Y_x + \Pi$, where $\Pi - \mathbb{E}(\Phi^T \Pi \Phi) = S$ for some $S \succ 0$. Then, since (1.19a) implies $Y_x - \mathbb{E}(\Phi^T Y_x \Phi) = L$, we have $\hat{Y}_x \succ Y_x$ and

$$\hat{Y}_x - \mathbb{E}(\Phi^T \hat{Y}_x \Phi) \succ L. \quad (1.41)$$

Also (1.23) gives $\operatorname{tr}(\Theta L) = \mathbb{E}[d^T Y_x d]$, and it follows that

$$\mathbb{E}(d^T \hat{Y}_x d) \ge \operatorname{tr}(\Theta L). \quad (1.42)$$

Using (1.41) and (1.42) it can be shown that the condition

$$\begin{bmatrix} \hat{Y}_x - L & 0 & 0 \\ 0 & \beta I & \hat{g} \\ 0 & \hat{g}^T & \gamma \end{bmatrix} - \mathbb{E}\left\{ \begin{bmatrix} \Phi^T \\ B^T \\ d^T \end{bmatrix} \hat{Y}_x \begin{bmatrix} \Phi & B & d \end{bmatrix} \right\} \succeq 0 \qquad (1.43)$$

is satisfied for $\gamma \geq \mathbb{E}(d^T \hat{Y}_x d) \geq \mathrm{tr}(\Theta L)$ and sufficiently large (but finite) β. Pre- and post-multiplying (1.43) by $\begin{bmatrix} x_k^T & f_{0|k}^T & 1 \end{bmatrix}$ and its transpose gives

$$x_k^T \hat{Y}_x x_k - \mathbb{E}_k(x_{k+1}^T \hat{Y}_x x_{k+1}) + \beta \|f_{0|k}^*\|^2 + 2\hat{g}^T f_{0|k}^* \leq x_k^T L x_k - \mathrm{tr}(\Theta L).$$

Finally, taking expectations and summing both sides of this inequality over $k = 0, \ldots n-1$ yields (1.40). □

It is now possible to derive bounds on the convergence of the closed loop system state by combining Lemma 1.10 and Theorem 1.11.

Theorem 1.12. *Under Algorithm 2, the state of (1.2) satisfies*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}_0(x_k^T L x_k) \leq \mathrm{tr}(\Theta L). \qquad (1.44)$$

Proof. If $\hat{g}$ in Theorem 1.11 is chosen as $\hat{g} = \beta g / \underline{\sigma}(D)$, with $\underline{\sigma}(\cdot)$ denoting the minimum singular value, we obtain

$$\mathbb{E}_0\big[\beta \|f_{0|k}^*\|^2 + 2\hat{g}^T f_{0|k}^*\big] \leq \frac{\beta}{\underline{\sigma}(D)} \mathbb{E}_0\big[\|f_{0|k}^*\|_D^2 + 2g^T f_{0|k}^*\big]$$

so that the combination of (1.38) and (1.40) yields

$$\sum_{k=0}^{n-1} \big[\mathbb{E}_0(x_k^T L x_k) - \mathrm{tr}(\Theta L)\big] \leq \frac{\beta}{\underline{\sigma}(D)} \big[V_{f,0}^* - \mathbb{E}_0(V_{f,n-1}^*)\big] + x_0^T \hat{Y}_x x_0 - \mathbb{E}_0(x_n^T \hat{Y}_x x_n)$$

To complete the proof, note that each term on the RHS of this inequality is finite since $V_{f,0}^*$ and $x_0^T \hat{Y}_x x_0$ are finite by assumption and since $V_{f,n-1}^*$ has a finite lower bound given by $-P_{1f} P_f^{-1} P_{f1}$. □

Theorem 1.12 allows the following comparison of Algorithm 2 with the linear feedback law $u_k = K x_k$. If $u_k = K x_k$, then $\mathbb{E}_0(\|x_k\|_Q^2 + \|u_k\|_R^2) = \mathbb{E}_0(x_k^T L x_k)$,

and since $\bar{x}_k \to 0$ as $k \to \infty$ along trajectories of the closed loop system under $u_k = K x_k$, from (1.8a) we obtain $\mathbb{E}_0(x_k^T L x_k) \to \mathrm{tr}(\Theta L)$ so the time-average of $\mathbb{E}_0(x_k^T L x_k)$ must converge to $\mathrm{tr}(\Theta L)$ as $k \to \infty$. However in (1.44), $\mathrm{tr}(\Theta L)$ is an upper bound for the time-average of the expected value of $x_k^T L x_k$ in closed loop operation under Algorithm 2. Thus Theorem 1.12 implies that this bound is achieved even though Algorithm 2 allows for the handling of constraints that are not accounted for explicitly by fixed linear feedback law $u_k = K x_k$.

1.6 Supermartingale convergence analysis

For $\kappa \geq 1$ it is possible to present an alternative treatment of the convergence analysis which allows some insight into the behaviour of the actual cost V, and this is undertaken in this section. The following lemma shows that the predicted cost corresponding to the solution of (1.35) or (1.36) necessarily decreases whenever $\|x_k\|_Q^2 + \|u_k\|_R^2$ is sufficiently large.

Lemma 1.13. *Let $V^*(x_k) = V(x_k, \mathbf{f}^*(x_k))$, then the receding horizon application of Algorithm 2 ensures that*

$$V^*(x_k) - \mathbb{E}_k\big[V^*(x_{k+1})\big] \geq \|x_k\|_Q^2 + \|u_k\|_R^2 - \kappa^2 \mathrm{tr}(\Theta L). \qquad (1.45)$$

Proof. From the expression for V in (1.17) we have

$$V^*(x_k) - \mathbb{E}_k\big[V(x_{k+1}, M\mathbf{f}^*(x_k))\big] = (1-\kappa^2) z_k^T \big[X - \mathbb{E}_k(\Psi^T X \Psi)\big] z_k + \kappa^2 z_k^T \big[Y - \mathbb{E}_k(\Psi^T Y \Psi)\big] z_k. \qquad (1.46)$$

However (1.18) and (1.19) imply that

$$z_k^T \big[X - \mathbb{E}_k(\Psi^T X \Psi)\big] z_k = \|x_k\|_Q^2 + \|u_k\|_R^2 - z_k^T \big[\mathbb{E}(\Psi^T X \Psi) - \bar{\Psi}^T X \bar{\Psi}\big] z_k \qquad (1.47)$$

$$z_k^T \big[Y - \mathbb{E}_k(\Psi^T Y \Psi)\big] z_k = \|x_k\|_Q^2 + \|u_k\|_R^2 - \mathrm{tr}(\Theta L) \qquad (1.48)$$

and the definitions of Ψ and z_k in (1.16a) give

$$z_k^T \big[\mathbb{E}(\Psi^T X \Psi) - \bar{\Psi}^T X \bar{\Psi}\big] z_k = \sum_{j=1}^{m} \|\Phi^{(j)} x_k + B^{(j)} E \mathbf{f}_k + d^{(j)}\|_{X_x}^2. \qquad (1.49)$$

Setting $\mathbf{f}_k = \mathbf{f}^*(x_k)$ in (1.49) and substituting (1.47-1.49) into (1.46) gives:

$$V^*(x_k) - \mathbb{E}_k\big[V(x_{k+1}, M\mathbf{f}^*(x_k))\big] = \|x_k\|_Q^2 + \|u_k\|_R^2 - \kappa^2 \mathrm{tr}(\Theta L) + (\kappa^2 - 1)\sum_{j=1}^{m} \|\Phi^{(j)} x_k + B^{(j)} E\mathbf{f}^*(x_k) + d^{(j)}\|_{X_x}^2 \quad (1.50)$$

However the objective in (1.35) and the constraint in (1.36) ensure that, for all realizations of model uncertainty at time k, the optimal value of the cost $V^*(x_{k+1})$ can be no greater than $V(x_{k+1}, M\mathbf{f}^*(x_k))$. Moreover, if $\kappa \geq 1$, then the last term on the RHS of (1.50) is positive, and hence this implies (1.45). □

Lemma 1.13 can be used to obtain an asymptotic bound on $\|x_k\|_Q^2 + \|u_k\|_R^2$ in closed-loop operation.

Theorem 1.14. *For $\kappa \geq 1$, under Algorithm 2 the closed-loop state and input trajectory satisfy*

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \mathbb{E}_0(\|x_k\|_Q^2 + \|u_k\|_R^2) \leq \kappa^2 \mathrm{tr}(\Theta L). \quad (1.51)$$

Proof. Taking expectations and summing both sides of the inequality (1.45) over $k = 0, \ldots, n-1$ gives

$$\sum_{k=0}^{n-1} \big[\mathbb{E}_0(\|x_k\|_Q^2 + \|u_k\|_R^2) - \kappa^2 \mathrm{tr}(\Theta L)\big] \leq V^*(x_0) - \mathbb{E}_0\big[V^*(x_n)\big] \quad (1.52)$$

which implies (1.51) since $V^*(x_0)$ is finite and, by (1.8a) and (1.12) $V^*(x_n)$ has a finite minimum value. □

For $\kappa = 1$, Theorem 1.14 implies that the stage cost $\|x_k\|_Q^2 + \|u_k\|_R^2$ in closed loop operation converges in mean to a value no greater than that obtained along the predicted trajectories of (1.15). However for $\kappa > 1$ the receding horizon controller places greater emphasis on minimizing the variance, measured by the sum of $\|x - \bar{x}\|_Q^2 + \|u - \bar{u}\|_R^2$, at the expense of the mean, measured by the sum of $\|\bar{x}\|_Q^2 + \|\bar{u}\|_R^2$, and the asymptotic bounds on the mean of $\|x\|_Q^2 + \|u\|_R^2$ consequently increase.

If the additive disturbance d_k were non-stationary with $\lim_{k\to\infty} \mathbb{E}(d_k d_k^T) = 0$, so that $\Theta = 0$ in (1.8b), then (1.45) would imply mean square stability of the closed loop system and hence ensure convergence: $x_k \to 0$ with probability 1 (w.p.1) as $k \to \infty$ Kushner (1971). For the more general problem considered here, mean square stability does not apply, however it is possible to obtain an alternative characterization of stability based on convergence of x_k to a set Ω defined by

$$\Omega = \big\{v : v^T Q v \leq \kappa^2 \mathrm{tr}(\Theta L)\big\}. \quad (1.53)$$

Theorem 1.15. *If $x_0 \notin \Omega$, then the closed loop system under Algorithm 2 satisfies $x_k \in \Omega$ for some $k > 0$ w.p.1.*

Proof. Define the sequence $\{\hat{x}_k,\ k = 1, 2, \ldots\}$ by

$$\hat{x}_k = \begin{cases} x_k & \text{if } \|x_i\|_Q^2 > \kappa^2 \mathrm{tr}(\Theta L) \quad \text{for all } i = 0, \ldots k-1 \\ \hat{x}_{k-1} & \text{if } \|x_i\|_Q^2 \leq \kappa^2 \mathrm{tr}\ (\Theta L) \quad \text{for any } i = 0, \ldots, k-1 \end{cases} \quad (1.54)$$

Then (1.45) implies

$$V^*(\hat{x}_k) - \mathbb{E}_k\big[V^*(\hat{x}_{k+1})\big] \geq \|\hat{x}_k\|_Q^2 - \kappa^2 \mathrm{tr}(\Theta L) \geq 0 \quad (1.55)$$

for all k, which establishes that $\{V^*(\hat{x}_k),\ k = 0, 1, \ldots\}$ is a supermartingale. Since $V^*(\hat{x}_k)$ is lower-bounded, it follows that $\|\hat{x}_k\|_Q^2 \to \kappa^2 \mathrm{tr}(\Theta L)$ w.p.1. (Kushner, 1971). □

Theorem 1.15 demonstrates that every state trajectory of the closed loop system (1.2) converges to the set Ω, although subsequently it may not remain in Ω. The same result shows that the state continually returns to Ω.

1.7 Numerical Example

This section demonstrates the action of the proposed stochastic predictive control law by applying it to a randomly selected, open loop unstable plant

with model parameters:

$$\bar{A} = \begin{bmatrix} 1.00 & 0.16 & -0.50 \\ -0.06 & 0.35 & -0.18 \\ -0.52 & 0.10 & 0.74 \end{bmatrix}, \quad \bar{B} = \begin{bmatrix} 0 \\ -0.28 \\ -0.12 \end{bmatrix}, \quad C^T = \begin{bmatrix} 0.11 \\ -1.86 \\ 0.71 \end{bmatrix}$$

$$A^{(1)} = \begin{bmatrix} -0.09 & -0.05 & 0.07 \\ -0.09 & 0 & 0.02 \\ 0.06 & 0 & 0.07 \end{bmatrix}, \quad B^{(1)} = \begin{bmatrix} 0.01 \\ -0.02 \\ 0.02 \end{bmatrix}, \quad d^{(1)} = \begin{bmatrix} 0.03 \\ 0.04 \\ 0.09 \end{bmatrix}$$

$$A^{(2)} = \begin{bmatrix} -0.06 & 0.07 & 0.06 \\ -0.03 & 0.03 & 0.09 \\ 0.10 & 0.07 & 0.05 \end{bmatrix}, \quad B^{(2)} = \begin{bmatrix} 0.02 \\ -0.09 \\ 0.01 \end{bmatrix}, \quad d^{(2)} = \begin{bmatrix} 0.07 \\ 0.10 \\ 0.09 \end{bmatrix}$$

The reference is $r = 10$, which gives $x_{ss} = \begin{bmatrix} -4.1466 & -6.4038 & -2.0492 \end{bmatrix}^T$ and $u_{ss} = 17.1$. The linear feedback gain K is chosen to be the unconstrained optimal computed using Algorithm 1. Figure 2 shows the evolution of the eigenvalues of ${P_x}^{(k)}$ at iteration k of Algorithm 1 for $\kappa = 1.9$ (for which $K = \begin{bmatrix} -11.7 & -0.913 & 9.02 \end{bmatrix}$) and for $\kappa = 0.5$ ($K = \begin{bmatrix} -10.1 & -0.750 & 7.87 \end{bmatrix}$).

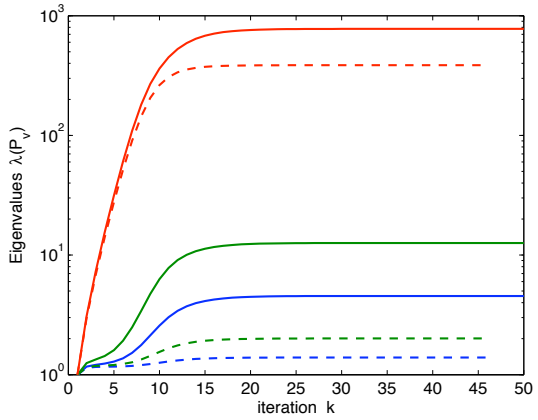

Figure 2: Variation of eigenvalues of P_x with iteration of Algorithm 1, for $\kappa = 1.9$ (solid lines) and $\kappa = 0.5$ (dashed lines)

For initial condition $x_0 = \begin{bmatrix} 15.9 & -6.4 & -2.05 \end{bmatrix}^T$, Fig. 3 shows the ensemble

average of the system output response y_k under receding horizon predictive control with $\kappa = 1.9$. Note from this plot that the output mean converges to a lower value than r, this being due to the trade off in the cost between mean and variance. For comparison, the output response under the linear feedback law is also shown, and the dashed lines show the bounds of mean ± 1 standard deviation. Clearly, linear feedback has much greater variance but achieves zero mean error in steady state. The horizontal black lines indicate $r \pm 1$ standard deviation as given by $(C\Theta C^T)^{0.5}$.

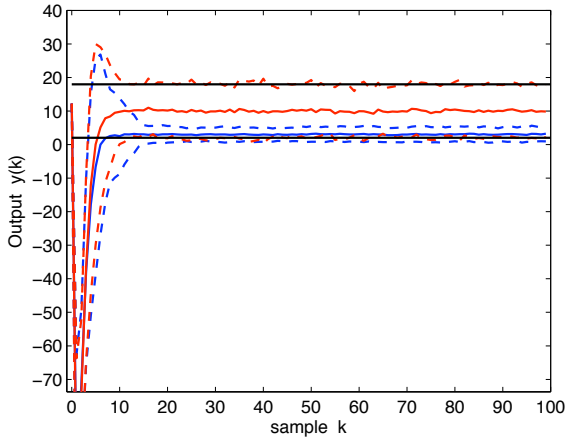

Figure 3: For $\kappa = 1.9$: ensemble averages (solid lines) and mean ± 1 standard deviation (dashed) for output responses $y(k)$ under Algorithm 2 (blue) and linear feedback (red); bounds $r \pm \sqrt{C\Theta_{ss}C^T}$ shown in black.

If constraints are inactive, the difference between output responses from receding horizon and linear feedback laws becomes less significant as κ is reduced (Fig. 4). However the variance of the receding horizon controller is again reduced when constraints are active. This is shown in Fig. 5, where the constraint that the expected value of y_k should be less than 15 is imposed.

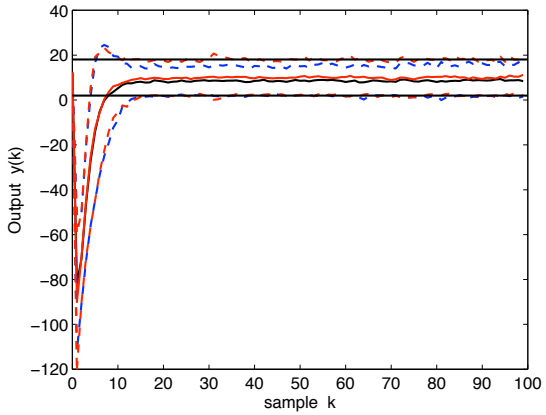

Figure 4: For $\kappa = 0.5$: ensemble averages (solid lines) and mean ± 1 standard deviation (dashed) for $y(k)$ under receding horizon control (blue) and linear feedback (red).

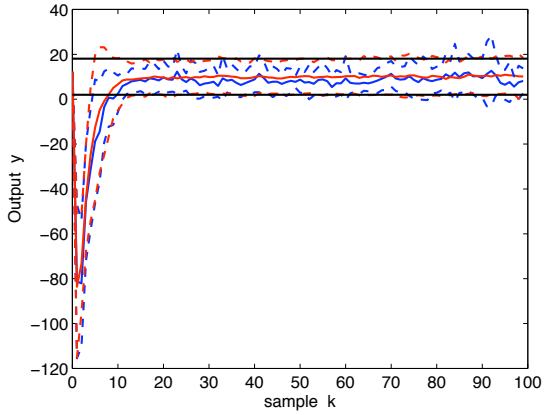

Figure 5: For $\kappa = 0.5$ and constraint $\mathbb{E}(y) \leq 15$: ensemble averages (solid) and bounds showing mean $\pm$ one variance (dashed) for $y(k)$ under Algorithm 2 (blue) and linear feedback (red).

2 Probabilistic Constraints

The constraints handled by predictive control strategies are traditionally treated either as hard constraints in the sense that they may never be violated, or as soft constraints, in which case the degree to which they are violated must be minimized in some sense. The probabilistic constraints introduced in this section here are a form of soft input/state constraint, the probability of violation of which is subject to hard limits. This form of probabilistic constraint has the advantage that it can take account of the distribution of model or measurement uncertainty, and thus avoid the conservativeness that can be introduced by using a hard-constraint strategy based on the worst-case uncertainty, which may be highly unlikely. Another important advantage is that probabilistic constraints are easier to design and interpret than soft-constraint strategies that are based on incorporating penalties on constraint violation into the MPC cost.

This section considers the same state space model and uncertainty description as used in Section 1. We also make use of the autonomous prediction model (1.16a) as well as the framework for defining the MPC cost and receding horizon optimization problem, and thus inherit the stability and convergence properties discussed in Section 1.5. The probabilistic constraints are defined in respect of a system output ψ_k, which may be subject to additive and multiplicative uncertainty:

$$\begin{aligned} \psi_k &= C_k x_k + D_k u_k + \eta_k, \quad \psi_k \in \mathbb{R}^{n_\psi} \\ \begin{bmatrix} C_k & D_k & \eta_k \end{bmatrix} &= \begin{bmatrix} \bar{C} & \bar{D} & 0 \end{bmatrix} + \sum_{j=1}^{m} \begin{bmatrix} C^{(j)} & D^{(j)} & \eta^{(j)} \end{bmatrix} q_k^{(j)} \end{aligned} \tag{2.1}$$

with $q_k = \begin{bmatrix} q_k^{(1)} & \cdots & q_k^{(m)} \end{bmatrix}^T \sim \mathcal{N}(0, I)$.

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 $N_{\max}$:

$$\frac{1}{N_c} \sum_{i=0}^{N_c-1} \Pr\{\psi_{k+i} \notin I_\psi\} \leq N_{\max}/N_c. \tag{2.2}$$

This statement can be translated into probabilistic constraints on the model state (as discussed in Section 2.2), and hence into constraints invoked in the online MPC optimization (discussed in Section 2.3). Within this framework soft input constraints are a special case of (2.1-2.2) with $C_k = 0$, $D_k = I$, $\eta_k = 0$.

2.1 Propagation of uncertainty

A systematic, non-conservative and computationally 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). The difficulty is that methods based on applying probabilistic constraints to predictions require the receding horizon optimization to be performed over closed-loop feedback policies, since otherwise there is no guarantee of the probabilistic constraints begin satisfied in closed-loop operation. This has had the effects of either restricting stochastic MPC to limited uncertainty classes (such as the additive uncertainty considered in van Hessem et al. (2001)) or requiring enormous computational effort to solve an approximate stochastic dynamic programming problem online (e.g. Batina et al., 2002).

This section describes a computationally convenient approach proposed in Cannon et al. (2008b,c), which is based on the autonomous description (1.16a) of the dynamics governing the evolution of future input and state predictions. We use Markov chain models to approximate the closed-loop evolution of the probability distributions for predicted states over the prediction horizon. These models are based on a discretization of the state space, which is computed offline using conditions for probabilistic invariance to ensure specified bounds on transition probabilities within the Markov chain model. Thus the approach effectively derives deterministic constraints offline in order to ensure the expected rate of constraint violation in closed-loop operation remains within the permitted bounds.

2.2 Markov Chain models

This section describes a method of analysis that enables the conversion of soft constraints (2.2) into probabilistic constraints on the state of (1.2). 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 (1.3) 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: $\mathcal{E}_1 \subset \cdots \subset \mathcal{E}_r$ (Fig. 6 depicts the case of $r = 3$), but $r = 2$ is used here to simplify presentation.

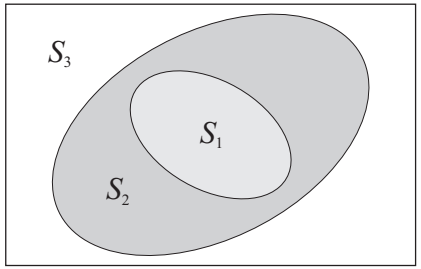

Figure 6: A pair of probabilistically invariant sets S_1 and S_2, and an invariant set S_3

Define the conditional probabilities

$$\Pr(\psi_k \notin I_\psi \mid x_k \in S_j) = p_j, \quad j = 1, 2 \tag{2.3}$$

where I_ψ is the constraint interval for ψ in (2.2). Under the assumption that p_1 is small, so that S_1 is a safe region of state space, it is convenient (though possibly conservative), to assume that S_2 is unsafe and thus assume $p_2 > p_1$. Define also the matrix of transition probabilities

$$\Pi = \begin{bmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{bmatrix} \tag{2.4}$$

where p_{ij} is the probability that the online algorithm steers the state in one

step from from S_j to S_i (Fig. 7). 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

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

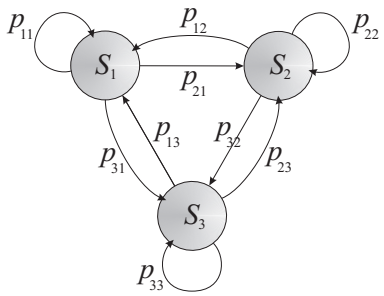

Figure 7: A 3-state Markov chain model with states S_1, S_2, S_3 and transition probabilities p_{ij}

By definition the transition matrix Π has the property that each column sums to 1. Its eigenvalue/vector decomposition therefore has the structure (see e.g. Kushner, 1971):

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

It follows that the rate at which constraint violations accumulate as $i \to \infty$ given $x_k \in S_j$ tends to

$$R_j = \begin{bmatrix} p_1 & p_2 \end{bmatrix} w_1 v_1^T e_j, \quad j = 1, 2 \tag{2.5}$$

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 (2.2), 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 (2.3),(2.4) ensures that the soft constraints (2.2) on ψ are satisfied.

The argument is illustrated by Figure 8. Here the minimum horizon over which the rate of accumulation of constraint violations satisfies the required bound is 8, and hence the constraint (2.2) is satisfied for this example if $N_c \geq 8$. It can also be seen from the upper solid line (corresponding to the expected number of constraint violations given initial conditions in S_2), that the expected rate of constraint violations may initially be higher than the maximum rate $\bar{R} = N_{\max}/N_c$. This is possible because the slope of the solid lines converge to the asymptotic rates $R_j < \bar{R}$, $j = 1, 2$.

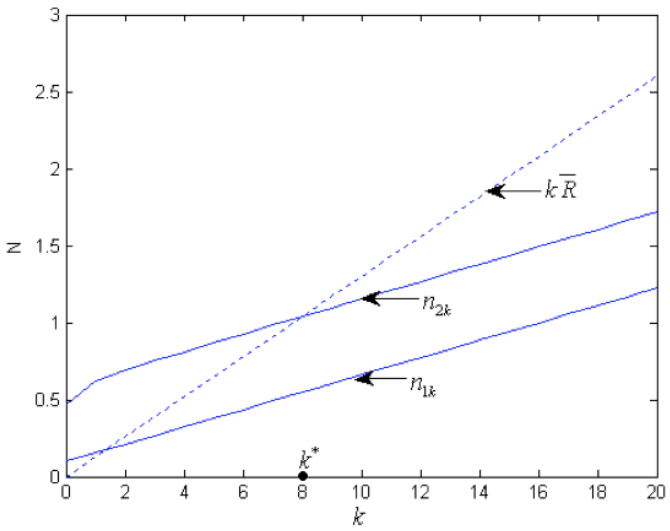

Figure 8: The expected number of constraint violations (N) over a k-step horizon

2.3 Probabilistically invariant sets

In the scenario discussed in section 2.2, 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 1 (Cannon et al., 2008b)**.** 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 defined in the state space of the prediction model (1.16a). For convenience the prediction model is restated here in an equivalent form with an explicit additive disturbance term:

$$\zeta_{k+i+1|k} = \Xi_{k+i}\zeta_{k+i|k}+\delta_{k+i}, \quad \zeta_{k|k} = \begin{bmatrix} x_k \\ \mathbf{f}_k \end{bmatrix}, \quad \delta_k = \begin{bmatrix} d_k \\ 0 \end{bmatrix}, \quad \Xi_k = \begin{bmatrix} \Phi_k & B_k E \\ 0 & M \end{bmatrix}. \tag{2.6}$$

We define ellipsoids $\mathcal{E} \subset \mathbb{R}^{n_x+Nn_u}$ and $\mathcal{E}_x \subset \mathbb{R}^{n_x}$ in the space of ζ and x respectively, by

$$\left.\begin{aligned} \mathcal{E} &= \{\zeta : \zeta^T \hat{P} \zeta \leq 1\} \\ \mathcal{E}_x &= \{x : x^T \hat{P}_x x \leq 1\} \end{aligned}\right\} \quad \hat{P}_x = (\Gamma^T \hat{P}^{-1} \Gamma)^{-1}$$

where $\Gamma^T = \begin{bmatrix} I & 0 \end{bmatrix}$ is the projection matrix such that $x_{k+i} = \Gamma^T \zeta_{k+i|k}$. With this definition it is easy to show that $\mathcal{E}_x$ is the projection of $\mathcal{E}$ onto the x-subspace (i.e. $\mathbf{f}$ exists such that $\zeta = [x^T \ \ \mathbf{f}^T]^T \in \mathcal{E}$ whenever $x \in \mathcal{E}_x$). Let $\mathcal{Q}$ denote a set that contains the vector of uncertain parameters in (1.3) with a specified probability p:

$$\Pr\{q_k \in \mathcal{Q}\} \geq p. \tag{2.7}$$

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., 2008b) used ellipsoidal confidence regions derived from this hypersphere to compute i.w.p. sets, but to accommodate the additive uncertainty in (1.2), 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 2.1. *$\mathcal{E}_x$ is i.w.p. p under (1.14) for any $\mathbf{f}_k$ such that $\zeta_{k|k} \in \mathcal{E}$ if there exists a scalar $\lambda \in [0, 1]$ satisfying*

$$\begin{bmatrix} \hat{P}_x^{-1} & \Gamma^T \Xi(q^i) \hat{P}^{-1} & \Gamma^T \gamma(q^i) \\ \hat{P}^{-1} \Xi(q^i)^T \Gamma & \lambda \hat{P}^{-1} & 0 \\ \gamma(q^i)^T \Gamma & 0 & 1-\lambda \end{bmatrix} \geq 0 \tag{2.8}$$

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

Proof. From (2.7) it follows that $\Pr(x_{k+1}^T \hat{P}_x x_{k+1} \leq 1) \geq p$ if

$$x_{k+1}^T \hat{P}_x x_{k+1} \leq 1 \quad \forall \zeta_{k|k} \in \mathcal{E}, \ \ \forall q_k \in \mathcal{Q} \tag{2.9}$$

where, under (1.14), x_{k+1} is given by $x_{k+1} = \Gamma^T \Xi(q_k) \zeta_{k|k} + \Gamma^T \gamma(q_k)$. By the $\mathcal{S}$-procedure (Boyd et al., 1994), (2.9) is equivalent to the existence of $\lambda \geq 0$ satisfying

$$1 - \Big(\Xi(q)\zeta + \gamma(q)\Big)^T \Gamma \hat{P}_x \Gamma^T \Big(\Xi(q)\zeta + \gamma(q)\Big) \geq \lambda(1 - \zeta^T \hat{P} \zeta)$$

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

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

Using Schur complements this can be expressed as a LMI in q, which, when invoked for all $q \in \mathcal{Q}$ is equivalent to (2.8) 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 $\zeta_{k|k}$ lies in $\mathcal{E}$. Re-writing (2.1) in the form:

$$\psi_k = \hat{C}(q_k)\zeta_{k|k} + \eta(q_k)$$
$$\hat{C}(q_k) = \bar{C}\Gamma^T + \bar{D}\hat{K} + \sum_{j=1}^m (C^{(j)}\Gamma^T + \tilde{D}^{(j)}\hat{K}) q_k^{(j)}$$

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

Lemma 2.2. $\Pr(\psi_k \notin I_\psi \mid \zeta_{k|k} \in \mathcal{E}) \leq 1 - p$ *if*

$$\psi_L \leq \eta(q^i) \leq \psi_U \tag{2.10}$$

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

$$\left[\hat{C}(q^i)\hat{P}^{-1}\hat{C}(q^i)^T\right]_{jj} \leq \left[\psi_U - \eta(q^i)\right]_j^2 \tag{2.11a}$$

$$\left[\hat{C}(q^i)\hat{P}^{-1}\hat{C}(q^i)^T\right]_{jj} \leq \left[\eta(q^i) - \psi_L\right]_j^2 \tag{2.11b}$$

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

Proof. For given q, $\max_{\zeta\in\mathcal{E}}[\hat{C}(q)\zeta]_j = [\hat{C}(q)\hat{P}^{-1}\hat{C}(q)^T]_{jj}^{1/2}$, and it follows from (2.7) that $\Pr(\psi_k \in I_\psi) \geq p$ whenever $\zeta_{k|k} \in \mathcal{E}$ if

$$\left[\hat{C}(q)\hat{P}^{-1}\hat{C}(q)^T\right]_{jj}^{1/2} \leq \left[\psi_U - \eta(q)\right]_j \tag{2.12a}$$

$$\left[\hat{C}(q)\hat{P}^{-1}\hat{C}(q)^T\right]_{jj}^{1/2} \leq \left[\eta(q) - \psi_L\right]_j \tag{2.12b}$$

for all $q \in \mathcal{Q}$ and $j = 1, \ldots, n_\psi$. Since (2.12a,b) are convex in q, the equivalent constraints (2.10),(2.11a,b) are obtained by invoking (2.12a,b) at each vertex of the polytope $\mathcal{Q}$. □

With $\mathcal{E}_1$ defined as $\mathcal{E}_x$ in the treatment outlined in section 2.2, the values of p_1 and p_{11} can specified using the constraints of Lemmas 2.1 and 2.2. To maximize the *safe* region of state space, it is clearly desirable to maximize $\mathcal{E}_x$, which can be formulated as

$$\underset{\hat{P}^{-1},\ \lambda\in[0,1]}{\text{maximize}}\ \det(\hat{P}_x^{-1}) \text{ subject to (2.8), (2.10) and (2.11a,b)} \tag{2.13}$$

If λ is a constant, then the constraints in (2.13) 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 (2.8), (2.10) and (2.11a,b), with the scalar λ fixed at a sequence of values in the interval $[0, 1]$.

In the case of input constraints: $u_L \leq u_k \leq u_U$, the constraints of lemma 2.2 reduce to $u_L \leq 0 \leq u_U$ and

$$\left[\hat{K}^T\hat{P}^{-1}\hat{K}\right]_{jj} \leq \left[u_U\right]_j^2, \quad \left[\hat{K}^T\hat{P}^{-1}\hat{K}\right]_{jj} \leq \left[u_L\right]_j^2 \tag{2.14}$$

for $j = 1, \ldots, n_u$. In the example of section 2.2 with $\mathcal{E}_1 = \mathcal{E}_x$, this implies $p_1 = 0$.

Corollary 2.3. *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 (2.2) will be satisfied in closed loop operation under algorithm 2.*

Proof. This follows directly from the assumptions on R_1, R_2 and the arguments of section 2.2. □

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 2.2 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 (2.8),(2.10),(2.11a,b) formulated, allowing $\hat{P}$ to be optimized by solving (2.13). 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 (2.2). 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.

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. For example, high-cycle fatigue damage constraints, which are considered in the example of Section 3.1, can be formulated using Miner's rule (Miner, 1945) and empirical models to combine the effects of stress cycles of differing amplitudes.

3 Example: Wind Turbine Blade Pitch Control

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.

3.1 System model and constraints

A simplified model of blade pitch rotation is given by

$$J\frac{d^2\beta}{dt^2} + c\frac{d\beta}{dt} = T_m - T_p \tag{3.1}$$

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 (3.1) 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:

$$y_{k+1} = a_{k,1}y_k + a_{k,0}y_{k-1} + b_{k,1}u_k + b_{k,0}u_{k-1} + w_k \tag{3.2}$$

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 (3.2) can be written in the form (1.2), 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

$$\begin{bmatrix} A_k & d_k \end{bmatrix} = \begin{bmatrix} \bar{A} & 0 \end{bmatrix} + \sum_{j=1}^{3} \begin{bmatrix} A^{(j)} & g^{(j)} \end{bmatrix} q_{k,j} \tag{3.3}$$

$$\bar{A} = \begin{bmatrix} 0 & -0.97 \\ 1 & 1.56 \end{bmatrix}, \; \bar{B} = \begin{bmatrix} -0.20 \\ -0.21 \end{bmatrix}, \; \begin{bmatrix} A^{(1)} & d^{(1)} \end{bmatrix} = \begin{bmatrix} 0 & -0.09 & 0 \\ 0 & 0.13 & 0.02 \end{bmatrix}$$

$$\begin{bmatrix} A^{(2)} & d^{(2)} \end{bmatrix} = \begin{bmatrix} 0 & 0.21 & 0 \\ 0 & -0.009 & -0.06 \end{bmatrix}, \; \begin{bmatrix} A^{(3)} & d^{(3)} \end{bmatrix} = \begin{bmatrix} 0 & -0.06 & 0 \\ 0 & 0.02 & 0.05 \end{bmatrix}$$

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

$$\begin{bmatrix} C_k & \eta_k \end{bmatrix} = \begin{bmatrix} \bar{C} & 0 \end{bmatrix} + \sum_{j=1}^{2} \begin{bmatrix} C^{(j)} & \eta^{(j)} \end{bmatrix} q_{k,j} \tag{3.4}$$

$$\bar{C} = \begin{bmatrix} 0 & 729 \end{bmatrix}, \; \bar{D} = 959$$

$$\begin{bmatrix} C^{(1)} & \eta^{(1)} \end{bmatrix} = \begin{bmatrix} 0 & 300 & 50 \end{bmatrix}, \; \begin{bmatrix} C^{(2)} & \eta^{(2)} \end{bmatrix} = \begin{bmatrix} 0 & 50 & 100 \end{bmatrix}$$

The number of degrees of freedom in predictions (1.14) 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 $p_{11}^0 = 0.9$, $p_{12}^0 = 0.8$, $N_{\max}/N_c =$

0.3, the permissible value for p_1 was found to be 0.2. For these values, the optimization (2.13) resulted in

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

The corresponding maximal area i.w.p. sets are shown in Figure 9.

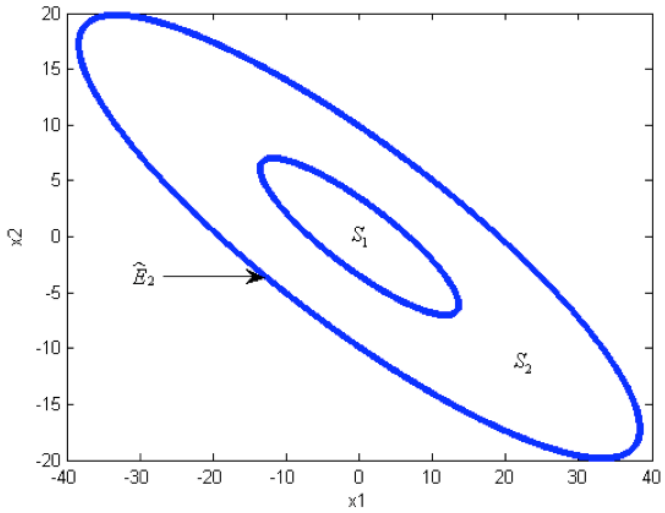

Figure 9: The sets $S_1 = \mathcal{E}_x$, and $S_2 = \mathcal{E}_{x2} - \mathcal{E}_{x1}$

3.2 Simulation results

Closed loop simulations of algorithm 2, 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_{12}^0, indicating that algorithm 2 satisfies the fatigue constraints.

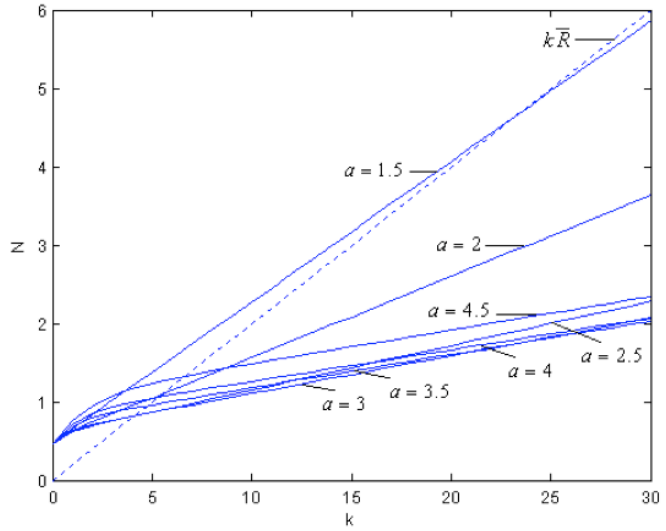

Figure 10: Expected number of constraint violations as p_{11} is varied

The design values of p_{11} and p_{12} can be optimized so as to minimize the value of i^* computed using the procedure of Section 2.2. This process is illustrated in Figures 10 and 11. Clearly the use of more sets S_i will allow further improvements in the value of i^*, since the accuracy of the Markov chain model in predicting constraint violations improves as a finer discretization of the state space is employed. The case of three sets S_1, S_2, S_3 is shown in Figure 12, and the corresponding rate of constraint violation is compared with the best achievable with two sets in Figure 13.

To establish the efficacy of algorithm 2, 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 2 gave an average closed loop cost of 257, whereas the average cost for $u_k = Kx_k$ was 325. Algorithm 2 achieves this improvement in performance by driving (during transients) the predictions hard against the

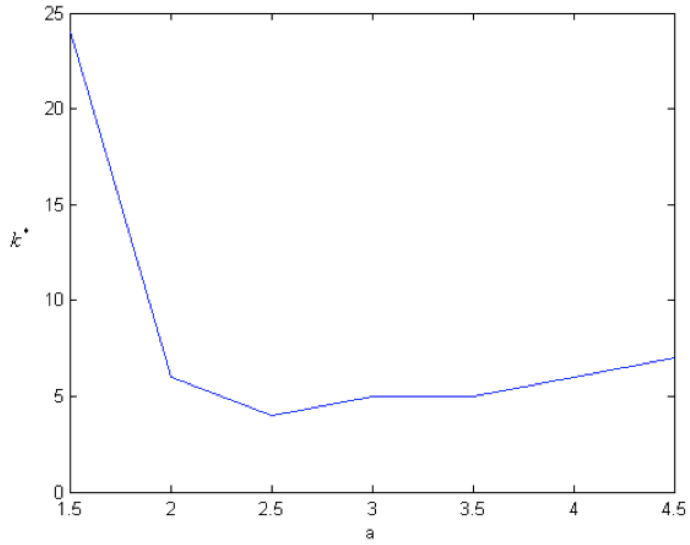

Figure 11: The variation of i^* with p_{11}

limits of the soft constraints. Both algorithm 2 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 $0 < n \leq 16$, indicate that $u_k = Kx_k$ exceeded the allowable limits during transients, whereas algorithm 2 gave average constraint violation rates less than $N_{\max}/N_c = 0.3$ for all $n \geq i^*$, where $i^* = 4$.

References

K.J. Astrom. *Introduction to Stochastic Control Theory*. Academic Press, New York, 1970.

I. Batina, A.A. Stoorvogel, and S. Weiland. Optimal control of linear, stochastic

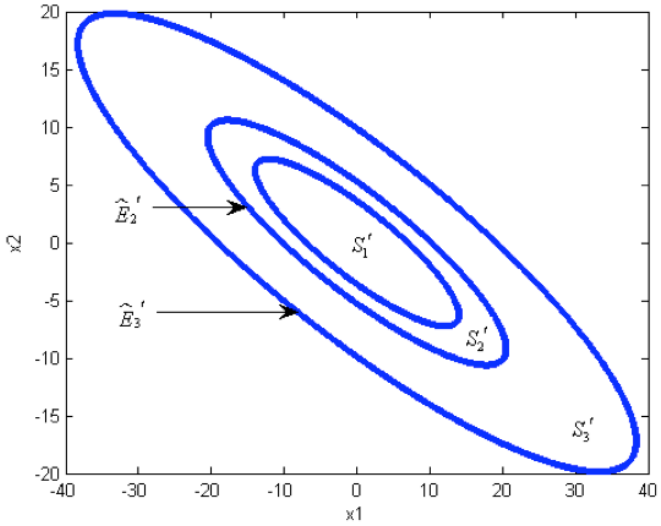

Figure 12: Sets S_1, S_2, S_3 for the case of $r = 3$

systems with state and input constraints. In *Proc. 41st IEEE Conf. Decision and Control*, pages 1564–1569, 2002.

R.E. Bellman. On the theory of dynamic programming. *Proceedings of the National Academy of Sciences*, 38:716–719, 1952.

G.E.P. Box and G.M. Jenkins. *Time series analysis: forecasting and control*. Holden-Day, San Francisco, 1976.

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

P.J. Brockwell and R.A. Davis. *Introduction to time series and forecasting*. Springer, New York, 2002.

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

M. Cannon, P. Couchman, and B. Kouvaritakis. MPC for stochastic systems.

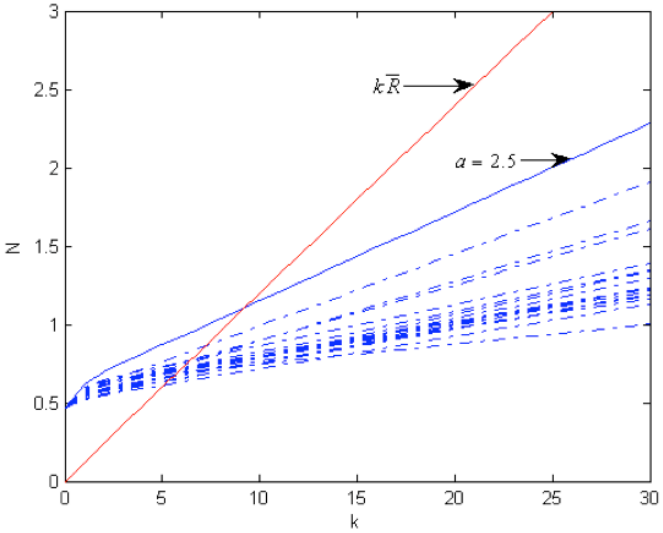

Figure 13: Expected numbers of constraint violations: for $r = 2$ with optimal p_{11} (solid) and for $r = 3$ for varying p_{11} and p_{22}

In *Assessment and Future Directions of Nonlinear Model Predictive Control*, volume 358 of *LNCIS*, pages 255–268. Springer, 2007.

M. Cannon, B. Kouvaritakis, and P. Couchman. Mean-variance receding horizon control for discrete time linear stochastic systems. In *Proc. 17th IFAC World Congress*, Seoul, 2008a.

M. Cannon, B. Kouvaritakis, and X. Wu. Model predictive control for systems with stochastic multiplicative uncertainty and probabilistic constraints. *Automatica*, 2008b. Accepted for publication.

M. Cannon, B. Kouvaritakis, and X. Wu. Probabilistic constrained mpc for multiplicative and additive stochastic uncertainty. In *Proc. 17th IFAC World Congress*, Seoul, 2008c.

P. Couchman, M. Cannon, and B. Kouvaritakis. Stochastic MPC with inequality stability constraints. *Automatica*, 42:2169–2174, 2006a.

P. Couchman, B. Kouvaritakis, and M. Cannon. MPC on state space models with stochastic input map. In *Proc. 45th IEEE Conf. Decision and Control*, pages 3216–3221, 2006b.

G. Freiling, S-R. Lee, and G. Jank. Coupled matrix Riccati equations in minimal cost variance control problems. *IEEE Trans. Automatic Control*, 44(3):556–560, 1999.

Z. Gajic and I. Borno. Lyapunov iterations for optimal control of jump linear systems at steady state. *IEEE Trans. Automatic Control*, 40(11):1971–1975, 1995.

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

T. Kailath. *Linear Systems*. Prentice-Hall, Englewood Cliffs, NJ, 1980.

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 P. Couchman. MPC as a tool for sustainable development integrated policy assessment. *IEEE Trans. Automatic Control*, 51(1):145–149, 2006.

E Kreyszig. *Advanced engineering mathematics*. Wiley, New York, 2006.

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

J.H. Lee and B.L. Cooley. Optimal feedback control strategies for state-space systems with stochastic parameters. *IEEE Trans. Automatic Control*, 43(10): 1469–1474, 1998.

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.

M.A. Miner. Cumulative damage in fatigue. *J. Appl. Mech.*, 12:A159–A164, 1945.

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. Dec. Contr.*, pages 4228–4233, 2001.

S.S. Zhu, D. Li, and S.Y. Wang. Risk control over bankruptcy in dynamic portfolio selection: A generalized mean-variance formulation. *IEEE Trans. Automatic Control*, 49(3):447–457, 2004.