MPC for Stochastic Systems

Mark Cannon, Paul Couchman, and Basil Kouvaritakis

Department of Engineering Science, University of Oxford, Oxford OX1 3PJ, UK
mark.cannon@eng.ox.ac.uk

Summary. Stochastic uncertainty is present in many control engineering problems, and is also present in a wider class of applications, such as finance and sustainable development. We propose a receding horizon strategy for systems with multiplicative stochastic uncertainty in the dynamic map between plant inputs and outputs. The cost and constraints are defined using probabilistic bounds. Terminal constraints are defined in a probabilistic framework, and guarantees of closed-loop convergence and recursive feasibility of the online optimization problem are obtained. The proposed strategy is compared with alternative problem formulations in simulation examples.

1 Introduction

The success of a Model Predictive Control (MPC) strategy depends critically on the choice of model. In most applications the plant model necessarily involves uncertainty, either endemic (e.g. due to exogenous disturbances) or introduced into the model to account for imprecisely known dynamics. It is usual in robust MPC to assume that uncertainty is bounded, or equivalently that it is random and uniformly distributed, and to adopt a worst case approach (e.g. [1, 2]). This is often considered to be overly pessimistic, even though it can be made less conservative through the use of closed-loop optimization [3, 4], albeit at considerable computational cost.

A more realistic approach, especially when uncertainty is known to be random but is not uniform, is to identify the distributions of uncertain model parameters and use these to solve a stochastic MPC problem. In many applications distributions for uncertain parameters can be quantified (e.g. as part of the model identification process), and some of the constraints are soft and probabilistic in nature (e.g. in sustainable development applications). Ignoring this information (by employing worst case performance indices and invoking constraints over all possible realizations of uncertainty) results in conservative MPC laws. This motivates the development of stochastic MPC formulations, which have been proposed for the case of additive disturbances (e.g. [5, 6, 7]) and for models incorporating multiplicative disturbances (e.g. [8, 9]).

Information on the distributions of stochastic parameters can be exploited in an optimal control problem by defining the performance index as the expected

R. Findeisen et al. (Eds.): Assessment and Future Directions, LNCIS 358, pp. 255–268, 2007.
springerlink.com

value of the usual quadratic cost. This approach is the basis of unconstrained LQG optimal control, and has more recently been proposed for receding horizon control [7, 10]. Both [10] and [7] consider input constraints, with [10] performing an open-loop optimization while [7] uses Monte Carlo simulation techniques to optimize over feedback control policies. This paper also considers constraints, but an alternative cost is developed based on bounds on predictions that are invoked with specified probability. The approach allows for a greater degree of control over the output variance, which is desirable for example in sustainable development, where parameter variations are large and the objective is to maximize the probability that the benefit associated with a decision policy exceeds a given aspiration level.

Probabilistic formulations of system constraints are also common in practice. For example an output may occasionally exceed a given threshold provided the probability of violation is within acceptable levels; this is the case for economic constraints in process control and fatigue constraints in electro-mechanical systems. Probabilistic constraints are incorporated in [11] through the use of statistical confidence ellipsoids, and also in [9], which reduces conservatism by applying linear probabilistic constraints directly to predictions without the need for ellipsoidal relaxations. The approach of [9] assumes Moving Average (MA) models with random coefficients, and achieves the guarantee of closed-loop stability through the use of an equality stability terminal constraints. The method is extended in [12] to more general linear models in which the uncertain parameters are contained in the output map of a state-space model, and to incorporate less restrictive inequality stability constraints.

The current paper considers the case of uncertain time-varying plant parameters represented as Gaussian random variables. This type of uncertainty is encountered for example in civil engineering applications (e.g. wind-turbine blade pitch control) and in financial engineering applications, where Gaussian disturbance models are common. Earlier work is extended in order to account for uncertainty in state predictions, considering in particular the definition of cost and terminal constraints to ensure closed-loop convergence and feasibility properties. For simplicity the model uncertainty is assumed to be restricted to Gaussian parameters in the input map, since this allows the distributions of predictions to be obtained in closed-form, however the design of cost and constraints extends to more general model uncertainty. After discussing the model formulation in section 2 and the stage cost in section 3, sections 4 and 5 propose a probabilistic form of invariance for the definition of terminal sets and define a suitable terminal penalty term for the MPC cost. Section 6 describes closed-loop convergence and feasibility properties, and the advantages over existing robust and stochastic MPC formulations are illustrated in section 7.

2 Multiplicative Uncertainty Class

In many control applications, stochastic systems with uncertain multiplicative parameters can be represented by MA models:

$$y_i(k) = \sum_{m=1}^{n_u} g_{im}^T(k)\tilde{u}_m(k-1), \quad \tilde{u}_m(k-1) = [u_m(k-n) \dots u_m(k-1)]^T \qquad (1)$$

where $u_m(k)$, $m = 1, \dots, n_u$, $y_i(k)$, $i = 1, \dots, n_y$, are input and output variables respectively, and the plant parameters $g_{im}(k)$ are Gaussian random variables. For convenience we consider the case of two outputs ($n_y = 2$): y_1 is taken to be *primary* (in that a probabilistic measure of performance on it is to be optimized) whereas y_2 is subject to probabilistic performance constraints and is referred to as *secondary.*

As a result of the linear dependence of the model (1) on uncertain plant parameters, the prediction of $y_i(k+j)$ made at time k (denoted $y_i(k+j|k)$) is normally distributed. Therefore bounds on $y_i(k+j|k)$ that are satisfied with a specified probability p can be formulated as convex (second-order conic) constraints on the predicted future input sequence. Bounds of this kind are used in [9] to derive a probabilistic objective function and constraints for MPC. These are combined with a terminal constraint that forces predictions to reach a pre-computed steady-state at the end of an N-step prediction horizon to define a stable receding horizon control law. Subsequent work has applied this methodology to a sustainable development problem using linear time-varying MA models [13].

Though often convenient in practice, MA models are non-parsimonious, and an alternative considered in [12] is given by the state space model:

$$x(k+1) = Ax(k) + Bu(k), \qquad y_i(k) = c_i^T(k)\, x(k), \quad i = 1, 2 \qquad (2)$$

where $x(k) \in \mathbb{R}^n$ is the state (assumed to be measured at time k), $u(k) \in \mathbb{R}^{n_u}$ is the input, and A, B are known constant matrices. The output maps $c_i(k) \in \mathbb{R}^n$, $i = 1, 2$ are assumed to be normally distributed: $c_i(k) \sim \mathcal{N}(\bar{c}_i, \Theta_{c,i})$, with $\{c_i(k), c_i(j)\}$ independent for $k \neq j$. The stability constraints of [9] are relaxed in [12], which employs less restrictive inequality constraints on the N step-ahead predicted state.

This paper considers a generalization of the model class in order to handle the case that the future plant state is a random variable. For simplicity we restrict attention to the case of uncertainty in the input map:

$$x(k+1) = Ax(k) + B(k)u(k), \;\; B(k) = \bar{B} + \sum_{r=1}^{L} q_r(k)B_r, \;\; y_i(k) = c_i^T x(k), \; i = 1, 2 \qquad (3)$$

where A, $\bar{B}$, B_i, c_i are known and $q(k) = [q_1(k) \;\cdots\; q_L(k)]^T$ are Gaussian parameters. We assume that $q(k) \sim \mathcal{N}(0, I)$ since it is always possible to define the model realization (A, B, C) so that the elements of $q(k)$ are uncorrelated, and that $\{q(k), q(j)\}$ are independent for $k \neq j$. Correlation between model parameters at different times could be handled by the paper's approach, but the latter assumption simplifies the expressions for the predicted covariances in section 5 below. The state $x(k)$ is assumed to be measured at time k. The paper focuses on the design of the MPC cost and terminal constraints so as to ensure closed-loop stability (for the case of soft constraints) and recursive feasibility with a pre-specified confidence level.

3 Performance Index and Constraints

The control objective is to regulate the expected value and variance of the primary output while respecting constraints on inputs and secondary outputs. We define the receding horizon cost function to be minimized online at time k as

$$J = \sum_{j=0}^{N-1} l(k+j|k) + L(k+N|k) \tag{4}$$

where

$$l(k+j|k) = \bar{y}_1^2(k+j|k) + \kappa_1^2\sigma_1^2(k+j|k) \tag{5}$$

with $\bar{y}_1(k+j|k) = \mathbb{E}_k y_1(k+j|k)$ and $\sigma_1^2(k+j|k) = \mathbb{E}_k\big[y_1(k+j|k) - \bar{y}_1(k+j|k)\big]^2$ denoting the mean and variance of $y_1(k+j|k)$ given the measurement $x(k)$ (we denote the expectation of a variable z given the measurement $x(k)$ as $\mathbb{E}_k z$).

This form of stage cost is used in preference to the more usual expectation MPC cost (e.g. [7, 10, 11]) because it enables the relative weighting of mean and variance to be controlled directly via the parameter κ_1, which can be interpreted in terms of probabilistic bounds on the prediction $y_1(k+j|k)$. To see this, let t_{lower} and t_{upper} be lower and upper bounds on $y_1(k+j|k)$ with a given probability p_1:

$$\begin{aligned} \Pr\big(y_1(k+j|k) \geq t_{\text{lower}}(k+j|k)\big) \geq p_1 \\ \Pr\big(y_1(k+j|k) \leq t_{\text{upper}}(k+j|k)\big) \geq p_1 \end{aligned} \tag{6}$$

then, since the predictions generated by (3) are normally distributed, it is easy to show that the stage cost (5) is equivalent to

$$l(k+j|k) = \tfrac{1}{2}t_{\text{lower}}^2(k+j|k) + \tfrac{1}{2}t_{\text{upper}}^2(k+j|k)$$

provided κ_1 satisfies $\mathfrak{N}(\kappa_1) = p_1$, where $\mathfrak{N}$ is the normal distribution function: $\Pr(z \leq Z) = \mathfrak{N}(Z)$ for $z \sim \mathcal{N}(0,1)$.

An important property of the stage cost is that it allows closed-loop stability under the MPC law to be determined by considering the optimal value of J as a stochastic Lyapunov function. The analysis (which is summarized in Section 6 below) is based on the following result.

Lemma 1. *If $\kappa_1 \geq 1$, then for any given input sequence $\{u(k), u(k+1), \ldots, u(k+j-1)\}$, the expectation of $l(k+j|k+1)$ conditional on time k satisfies:*

$$\mathbb{E}_k l(k+j|k+1) \leq l(k+j|k). \tag{7}$$

Proof. Re-writing (5) as $l(k+j|k) = \mathbb{E}_k y_1^2(k+j|k) + (\kappa_1^2 - 1)\sigma_1^2(k+j|k)$, and noting that

$$\begin{aligned} \mathbb{E}_k\big(\mathbb{E}_{k+1} y_1^2(k+j|k+1)\big) &= \mathbb{E}_k y_1^2(k+j|k) \\ \mathbb{E}_k \sigma_1^2(k+j|k+1) &= \sigma_1^2(k+j|k+1) \end{aligned}$$

we have

$$\mathbb{E}_k l(k+j|k+1) = l(k+j|k) - (\kappa_1^2 - 1)\big(\sigma_1^2(k+j|k) - \sigma_1^2(k+j|k+1)\big).$$

The required bound therefore holds if $\kappa_1 \geq 1$ since $\sigma_1^2(k+j|k) \geq \sigma_1^2(k+j|k+1)$. ∎

Remark 1. In accordance with Lemma 1 it is assumed below that $\kappa_1 \geq 1$, or equivalently that the bounds (6) are invoked with probability $p_1 \geq 84.1\%$ (to 3 s.f.). With $\kappa_1 = 1$, this formulation recovers the conventional expectation cost: $l(k+j|k) = \mathbb{E}_k y_1^2(k+j|k)$ for regulation problems.

Consider next the definition of constraints. Since output predictions are Gaussian random variables, we consider probabilistic (as opposed to hard) constraints:

$$\Pr\big(y_2(k+j|k) \leq Y_2\big) \geq p_2 \tag{8}$$

where Y_2 is a constraint threshold. Input constraints are assumed to have the form:

$$|u(k+j|k)| \leq U \tag{9}$$

where $u(k+j|k)$ is the predicted value of $u(k+j)$ at time k.

4 Terminal Constraint Set

Following the conventional dual mode prediction paradigm [14], predicted input trajectories are switched to a linear terminal control law: $u(k+j|k) = Kx(k+j|k)$, $j \geq N$ after an initial N-step prediction horizon. For the case of uncertainty in the output map (2), an ellipsoidal terminal constraint can be computed by formulating conditions for invariance and satisfaction of constraints (8),(9) under the terminal control law as LMIs [12]. However, in the case of the model (3), the uncertainty in the predicted state trajectory requires that a probabilistic invariance property is used in place of the usual deterministic definition of invariance when defining a terminal constraint set. We therefore impose the terminal constraint that $x(k+N|k)$ lie in a terminal set Ω with a given probability, where Ω is designed so that the probability of remaining within Ω under the closed-loop dynamics $x(k+1) = \Phi(k)x(k)$ is at least p_Ω, i.e.

$$\Pr(\Phi x \in \Omega) \geq p_\Omega \quad \forall x \in \Omega. \tag{10}$$

If constraints on the input and secondary output are satisfied everywhere within Ω, then this approach can be used to define a receding horizon optimization which is feasible with a specified probability at time $k+1$ if it is feasible at time k. Given that the uncertain parameters of (3) are not assumed bounded, this is arguably the strongest form of recursive feasibility attainable.

For computational convenience we consider polytopic terminal sets defined by $\Omega = \{x : v_i^T x \leq 1,\ i = 1,\ldots,m\}$. Denote the closed-loop dynamics of (3) under $u = Kx$ as

$$x(k+1) = \Phi(k)x(k), \quad \Phi(k) = \bar{\Phi} + \sum_{i=1}^{L} q_i(k)\Phi_i, \quad q(k) \sim \mathcal{N}(0, I) \tag{11}$$

(where $\bar{\Phi} = A + \bar{B}K$ and $\Phi_i = B_i K$), then confidence ellipsoids for q can be used to determine conditions on the vertices x_j, $j = 1, \ldots, M$ of Ω so that Ω is invariant with a given probability. Specifically, the condition $v_i^T \Phi x_j \leq 1$ is equivalent to $x_j^T [\Phi_1^T v_i \ \cdots \ \Phi_L^T v_i] q \leq 1 - x_j^T \bar{\Phi}^T v_i$, and, since $\|q\|^2$ is distributed as χ^2 with L degrees of freedom, it follows that $v_i^T \Phi x_j \leq 1$ with probability p_Ω if

$$r_\Omega \left\| x_j^T \left[\Phi_1^T v_i \cdots \Phi_L^T v_i\right] \right\|_2 \leq 1 - x_j^T \bar{\Phi} v_i \tag{12}$$

where r_Ω satisfies $\Pr(\chi^2(L) < r_\Omega^2) = p_\Omega$.

Lemma 2. *Ω is invariant under (11) with probability p_Ω, i.e.*

$$\Pr(v_i^T \Phi x \leq 1, \ i = 1, \ldots, m) \geq p_\Omega, \quad \forall x \in \Omega \tag{13}$$

if (12) is satisfied for $i = 1, \ldots, m$ and $j = 1, \ldots, M$.

Proof. If (12) holds for given j and $i = 1, \ldots, m$, then x_j necessarily satisfies $\Pr(\Phi x_j \in \Omega) \geq p_\Omega$ (since $\|q\|_2 \leq r_\Omega$ with probability p_Ω). Furthermore, invoking this condition for each vertex x_j implies (13), since (12) is convex in x_j. ■

The problem of maximizing Ω subject to (10) and the conditions that input constraints (9) and the secondary output constraint $y_2 \leq Y_2$ are met everywhere within the terminal set can be summarized as:

$$\begin{aligned} \text{maximize} \quad & \text{vol}(\Omega) \\ \text{subject to} \quad & r_\Omega \left\| x_j^T \left[\Phi_1^T v_i \cdots \Phi_L^T v_i\right] \right\|_2 \leq 1 - x_j^T \bar{\Phi} v_i, \\ & |K x_j| \leq U \\ & c_2 x_j \leq Y_2 \end{aligned} \tag{14}$$

in variables $\{v_i, i = 1, \ldots, m\}$ and $\{x_j, j = 1, \ldots, M\}$. This is a nonconvex problem, but for fixed $\{v_i\}$ the constraints are convex in $\{x_j\}$, enabling a sequence of one-step sets of increasing volume to be computed via convex programming. Therefore a (locally) optimal point for (14) can be found using a sequential approach similar to that of [15]. Furthermore, if Ω is defined as a symmetric low-complexity polytope (i.e. $\Omega = \{x : \|Wx\|_\infty \leq 1\}$, for full-rank $W \in \mathbb{R}^{n \times n}$), then the linear feedback gain K can be optimized simultaneously with Ω by including the vertex controls, u_j, $j = 1, \ldots, n$ as additional optimization variables in (14), where $u_j = K x_j$.

5 Terminal Penalty

To allow a guarantee of closed-loop stability, we define the terminal penalty in (4) as the cost-to-go over all $j \geq N$ under the terminal control law $u(k+j|k) = Kx(k+j|k)$. This section derives the required function $L(k+N|k)$ as a quadratic form based on the solution of a pair of Lyapunov equations, and shows that the Lyapunov-like property:

$$\mathbb{E}_k\big[L(k+N+1|k+1) + l(k+N|k+1)\big] \leq L(k+N|k) \tag{15}$$

holds whenever predictions at time $k+1$ are generated by the sequence

$$\mathbf{u}(k+1) = \{u(k+1|k), \dots, u(k+N-1|k), Kx(k+N|k+1)\} \tag{16}$$

where $\mathbf{u}(k) = \{u(k|k), u(k+1|k), \dots, u(k+N-1|k)\}$ is the predicted input sequence at time k and $u(k) = u(k|k)$.

To simplify notation, let $x_\delta = x - \bar{x}$, where $\bar{x}(k+j|k) = \mathbb{E}_k x(k+j|k)$, and define

$$Z_1(k+j|k) = \mathbb{E}_k\big[x(k+j|k)x^T(k+j|k)\big],$$
$$Z_2(k+j|k) = \mathbb{E}_k\big[x_\delta(k+j|k)x_\delta^T(k+j|k)\big].$$

Lemma 3. *If the terminal penalty in (4) is defined by*

$$L(k+N|k) = \mathrm{Tr}\big(Z_1(k+N|k)S_1\big) + (\kappa_1^2 - 1)\mathrm{Tr}\big(Z_2(k+N|k)S_2\big) \tag{17}$$

where $S_1 = S_1^T \succ 0$ and $S_2 = S_2^T \succ 0$ are the solutions of the Lyapunov equations

$$\bar{\Phi}^T S_2 \bar{\Phi} + c_1 c_1^T = S_2 \tag{18a}$$

$$\bar{\Phi}^T S_1 \bar{\Phi} + \sum_{i=1}^{L} \Phi_i^T \big(S_1 + (\kappa_1^2 - 1)S_2\big)\Phi_i + c_1 c_1^T = S_1 \tag{18b}$$

then $L(k+N|k)$ is the cost-to-go: $L(k+N|k) = \sum_{j=N}^{\infty} l(k+j|k)$ for the closed-loop system formed by (3) under the terminal control law $u(k+j|k) = Kx(k+j|k)$.

Proof. With $u(k+j|k) = Kx(k+j|k)$, it is easy to show that, for all $j \geq N$:

$$Z_1(k+j+1|k) = \bar{\Phi} Z_1(k+j|k)\bar{\Phi}^T + \sum_{i=1}^{L} \Phi_i Z_1(k+j|k)\Phi_i^T \tag{19a}$$

$$Z_2(k+j+1|k) = \bar{\Phi} Z_2(k+j|k)\bar{\Phi}^T + \sum_{i=1}^{L} \Phi_i Z_2(k+j|k)\Phi_i^T. \tag{19b}$$

Using these expressions and (18a,b) to evaluate $L(k+j+1|k)$, we obtain

$$L(k+j+1|k) + l(k+j|k) = L(k+j|k), \tag{20}$$

which can be summed over all $j \geq N$ to give

$$L(k+N|k) - \lim_{j\to\infty} L(k+j|k) = \sum_{j=N}^{\infty} l(k+j|k),$$

but $x(k+1) = \Phi x(k)$ is necessarily mean-square stable [16] in order that there exist positive definite solutions to (18a,b), and it follows that $\lim_{j\to\infty} L(k+j|k) = 0$. ∎

Remark 2. For any $j \geq 1$, Z_1 and Z_2 can be computed using

$$Z_2(k+j|k) = \sum_{i=0}^{j-1} \Psi_i \Psi_i^T, \quad \Psi_i = A^{j-1-i}\big[B_1 u(k+i|k) \cdots B_L u(k+i|k)\big] \tag{21a}$$

$$Z_1(k+j|k) = \bar{x}(k+j|k)\bar{x}^T(k+j|k) + Z_2(k+j|k) \tag{21b}$$

Therefore $L(k+N|k)$ is a quadratic function of the predicted input sequence:

$$L(k+N|k) = \mathbf{u}^T(k) H \mathbf{u}(k) + 2g^T \mathbf{u}(k) + \gamma$$

where $\mathbf{u}(k) = [u^T(k|k) \cdots u^T(k+N-1|k)]^T$ and H, g are constants.

Theorem 1. *If $L(k+N|k)$ is given by (17) and $L(k+N+1|k+1), l(k+N+1|k+1)$ correspond to the predictions generated by the input sequence (16), then (15) is satisfied if $\kappa_1 \geq 1$.*

Proof. From (19), (18), and $u(k+N|k+1) = Kx(k+N|k+1)$ it follows that $L(k+N+1|k+1) + l(k+N|k+1) = L(k+N|k+1)$. Furthermore, from (21a,b) we have

$$\mathbb{E}_k Z_2(k+N|k+1) = Z_2(k+N|k+1), \quad \mathbb{E}_k Z_1(k+N|k+1) = Z_1(k+N|k).$$

Combining these results, the LHS of (15) can be written

$$\mathbb{E}_k\big[L(k+N+1|k+1) + l(k+N|k+1)\big] = \mathrm{Tr}\big(Z_1(k+N|k)S_1\big) + (\kappa_1^2 - 1)\mathrm{Tr}\big(Z_2(k+N|k+1)S_2\big)$$

and therefore (15) holds if $\kappa_1 \geq 1$ since (21a) implies $Z_2(k+N|k+1) \preceq Z_2(k+N|k)$. ■

6 MPC Strategy and Closed-Loop Properties

The stage cost, terminal cost and terminal constraints are combined in this section to construct a receding horizon strategy based on the online optimization:

$$\underset{\mathbf{u}(k)}{\text{minimize}} \;\; J_k = \sum_{j=0}^{N-1} l(k+j|k) + L(k+N|k) \tag{22a}$$

subject to the following constraints, invoked for $j = 1, \ldots, N-1$:

$$|u(k+j|k)| \leq U \tag{22b}$$

$$\Pr\big(y_2(k+j|k) \leq Y_2\big) \geq p_2 \tag{22c}$$

$$\Pr\big(x(k+N|k) \in \Omega\big) \geq p_2 \tag{22d}$$

The MPC law is defined as $u(k) = u^*(k|k)$, where $\mathbf{u}^*(k) = \{u^*(k|k), \ldots, u^*(k+N-1|k)\}$ is optimal at time k, computed on the basis of the measured $x(k)$. From the plant model and parameter distributions, the constraints (22c) on y_2 can be written

$$\kappa_2\big(c_2^T Z_2(k+j|k)c_2\big)^{1/2} \le Y_2 - c_2^T \bar{x}(k+j|k) \tag{23}$$

where κ_2 satisfies $\mathfrak{N}(\kappa_2) = p_2$. Similarly, making use of confidence ellipsoids for $q(k)$, the terminal constraint (22d) can be expressed

$$r_1\big(v_i^T Z_2(k+j|k)v_i\big)^{1/2} \le 1 - v_i^T \bar{x}(k+j|k), \quad i = 1,\ldots,m \tag{24}$$

where r_1 is defined by $\Pr\big(\chi^2(NL) \le r_1^2\big) = p_2$. It follows that (22) is convex, and has the form of a second-order cone program (SOCP), enabling solution via efficient algorithms [17]. The stability properties of the MPC law can be stated as follows.

Theorem 2. *Assume that (22) is feasible at all times $k = 0, 1, \ldots$. Then $y_1(k) \to 0$, and $\|x(k)\|_2$ converges to a finite limit with probability 1 if (A, c_1) is observable.*

Proof. From Lemmas 1 and 3, the cost, $\tilde{J}_{k+1}$, for the suboptimal sequence $\tilde{\mathbf{u}}(k+1) = \{u^*(k+1|k), \ldots, Kx^*(k+N|k)\}$ at time $k+1$ satisfies $\mathbb{E}_k \tilde{J}_{k+1} \le J_k^* - y_1^2(k)$, where J_k^* is the optimal value of (22a). After optimization at $k+1$ we have

$$\mathbb{E}_k J_{k+1}^* \le \mathbb{E}_k \tilde{J}_{k+1} - y_1^2(k) \le J_k^* - y_1^2(k) \tag{25}$$

It follows that J_k converges to a lower limit and $y_1(k) \to 0$ with probability 1 [18]. Furthermore the definitions of stage cost (5) and terminal penalty (17) imply that

$$J_k = \sum_{j=0}^{\infty} c_1^T \bar{x}(k+j|k)\bar{x}^T(k+j|k)c_1 + \kappa_1^2 c_1^T Z_2(k+j|k)c_1$$

and, since $\sum_{j=0}^{\infty} c_1^T Z_2(k+j|k)c_1$ is positive definite in $\mathbf{u}(k)$ if (A, c_1) is observable, it follows that J_k is positive definite in $x(k)$ if (A, c_1) is observable. Under this condition therefore, $\|x(k)\|_2$ converges to a finite limit with probability 1. ∎

Note that the derivation of (25) assumes a pre-stabilized prediction model; the same convergence property can otherwise be ensured by using a variable horizon N.

The constraints (22b-d) apply only to predicted trajectories at time k, and do not ensure feasibility of (22) at future times. For example, at time $k+1$, (23) requires

$$\kappa_2\big(c_2^T Z_2(k+j|k+1)c_2\big)^{1/2} \le Y_2 - c_2^T \bar{x}(k+j|k+1), \quad j = 1,\ldots,N$$

where $\bar{x}(k+j|k+1)$ is a Gaussian random variable at time k, with mean $\bar{x}(k+j|k)$ and variance $c_2^T A^{j-1} Z_2(k+1|k) A^{j-1^T} c_2$. Therefore (23) is feasible at $k+1$ with probability p_2 if

$$\kappa_2\big(c_2^T Z_2(k+j|k+1)c_2\big)^{1/2} + \kappa_2\big(c_2^T A^{j-1} Z_2(k+1|k) A^{j-1^T} c_2\big)^{1/2} \le Y_2 - c_2^T \bar{x}(k+j|k)$$

holds for $j = 1,\ldots,N$ at time k; this condition is necessarily more restrictive than (23) since $Z_2(k+j|k) = Z_2(k+j|k+1) + A^{j-1} Z_2(k+1|k) A^{j-1^T}$. In order to provide a recursive guarantee of feasibility we therefore include additional constraints in the online optimization, as summarized in the following result.

Theorem 3. *If (23) and (24) are replaced in the MPC online optimization (22) by*

$$\sum_{l=0}^{j-1} \kappa_2 \big(c_2^T A^{j-1-l} Z_2(k+l+1|k+l) A^{j-1-l^T} c_2\big)^{1/2} \le Y_2 - c_2^T \bar{x}(k+j|k) \quad (26a)$$

$$\sum_{l=0}^{j-2} \kappa_2 \big(v_i^T A^{j-1-l} Z_2(k+l+1|k+l) A^{j-1-l^T} v_i\big)^{1/2} + r_j \big(v_i^T Z_2(k+N|k+j-1) v_i\big)^{1/2} \le 1 - v_i^T \bar{x}(k+N|k) \quad (26b)$$

for $j = 2, \ldots, N$, where r_j is defined by $\Pr(\chi^2((N+1-j)L) \le r_j) = p_2/p_\Omega^{j-1}$, then feasibility of (22) at time k implies feasibility at time $k+1$ with probability p_2.

Proof. Condition (26a) ensures that: (i) $\Pr(y(k+j|k) \le Y_2) \ge p_2$ for $j = 1, \ldots, N$; (ii) $\Pr(y(k+j|k+1) \le Y_2) \ge p_2$, $j = 2, \ldots, N$, is feasible at $k+1$ with probability p_2; and (iii) the implied constraints are likewise feasible with probability p_2 when invoked at $k+1$. Here (iii) is achieved by requiring that the constraints $\Pr(y(k+l|k+j) \le Y_2) \ge p_2$ be feasible with probability p_2 when invoked at $k+j$, $j = 2, \ldots, N-1$. Condition (26b) ensures recursive feasibility of (22d) with probability p_2 through the constraint that $\Pr(x(k+N|k+j) \in \Omega) \ge p_2/p_\Omega^j$, $j = 0, \ldots, N-1$ (and hence also $\Pr(x(k+N+j|k+j) \in \Omega) \ge p_2$) should be feasible with probability p_2. ∎

Incorporating (26a,b) into the receding horizon optimization leads to a convex online optimization, which can be formulated as a SOCP. However (26) and the constraint that Ω should be invariant with probability $p_\Omega > p_2^{1/(N-1)} \ge p_2$ are more restrictive than (22c,d), implying a more cautious control law.

Remark 3. The method of computing terminal constraints and penalty terms described in sections 4 and 5 is unchanged in the case that A contains random (normally distributed) parameters. However in this case state predictions are not linear in the uncertain parameters, so that the online optimization (22) could no longer be formulated as a SOCP. Instead computationally intensive numerical optimization routines (such as the approach of [7]) would be required.

Remark 4. It is possible to extend the approach of sections 4 and 5 to nonlinear dynamics, for example using linear difference inclusion (LDI) models. In the case that uncertainty is restricted to the linear output map, $y_j(k) = C_j(k)x(k)$ predictions then remain normally distributed, so that the online optimization, though nonconvex in the predicted input sequence, would retain some aspects of the computational convenience of (22).

7 Numerical Examples

This section uses two simulation examples to compare the stochastic MPC algorithm developed above with a generic robust MPC algorithm and the stochastic MPC approach of [12].

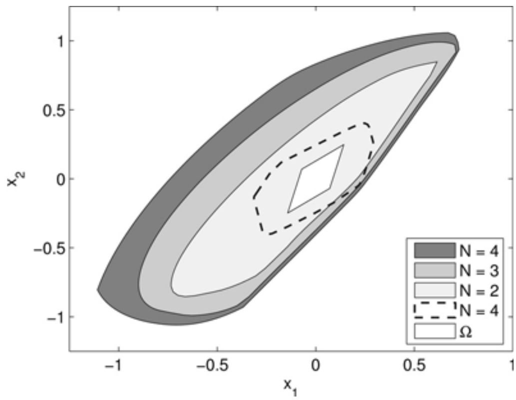

Fig. 1. Feasible initial condition sets for stochastic MPC ($p_2 = 0.85$) with varying N. Dashed line: feasible set of robust MPC based on 85% confidence level for $N = 4$.

First consider the plant model with

$$A = \begin{bmatrix} 1.04 & -0.62 \\ 0.62 & 1.04 \end{bmatrix} \; \bar{B} = \begin{bmatrix} 0 \\ 2 \end{bmatrix} \; B_1 = \begin{bmatrix} -0.12 \\ 0.02 \end{bmatrix} \; B_2 = \begin{bmatrix} 0.04 \\ -0.06 \end{bmatrix} \; c_1^T = \begin{bmatrix} 0 \\ -4.4 \end{bmatrix} \; c_2^T = \begin{bmatrix} 3 \\ 2.3 \end{bmatrix}$$

$U = 1$, $Y_2 = 1$, and $p_1 = p_2 = 0.85$. The offline computation for stochastic MPC involves maximizing a low-complexity polytopic set Ω subject to $\Pr(\Phi x \in \Omega) \geq p_2$ for all $x \in \Omega$; for this example the maximal Ω has an area of 0.055.

An alternative approach to MPC is to determine bounds on plant parameters corresponding to a confidence level of, say, p_2 by setting

$$B(k) = \bar{B} + \sum_{i=1}^{L} q_i(k) B_i, \quad |q(k)| \leq \mathfrak{N}^{-1}(p_2) \tag{27}$$

in (3), and then to implement a robust MPC law based on this approximate plant model. For a confidence level of $p_2 = 0.85$, the maximal low-complexity set Ω', which is robustly invariant for bounded parameter variations (27), is similar in size (area = 0.048) to Ω. The similarity is to be expected since the assumption of bounded uncertainty implies that the probability that $\Phi x \in \Omega'$ under the actual plant dynamics for any $x \in \Omega'$ is p_2.

A robust (min-max) MPC law employing open-loop predictions based on the parameter bounds of (27) is, however, significantly more conservative than the stochastic MPC law of (22) for the same confidence level. This can be seen in Fig. 1, which compares the feasible sets for the two control laws for $p_2 = 0.85$ (the feasible set for robust MPC decreases with increasing N for $N > 4$ since the plant is open-loop unstable). Closed-loop performance is also significantly worse: the closed-loop cost for robust MPC (based on the parameter bounds (27) with $p_2 = 0.85$), averaged over 10 initial conditions and 200 uncertainty realizations, is 63% greater than that for stochastic MPC (with $p_1 = p_2 = 0.85$). Figures 2 and 3 compare the closed-loop responses for a single initial condition and 20

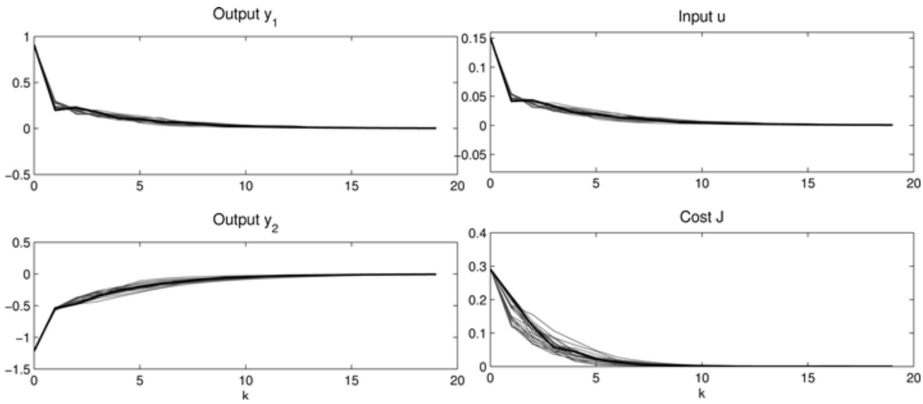

Fig. 2. Stochastic MPC closed-loop responses for $p_2 = 0.85$ and 20 uncertainty realizations (dark lines show responses for a single uncertainty realization)

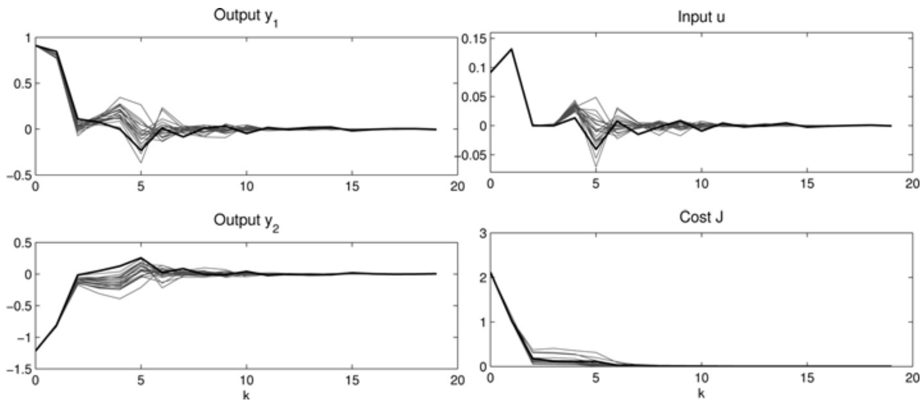

Fig. 3. Robust MPC closed-loop responses for 85% confidence levels and the same set of uncertainty realizations as in Fig. 2

uncertainty realizations. The higher degree of conservativeness and greater variability in Fig. 3 is a result of the robust min-max strategy, which attempts to control worst-case predictions based on the confidence bounds of (27), whereas the stochastic MPC strategy (Fig. 2) has direct control over the statistics of future predictions at each sampling instant.

Consider next the effects of approximating uncertainty in the input map as output map uncertainty. Modelling uncertainty in plant parameters as output map uncertainty simplifies MPC design since state predictions are then deterministic, but can result in a higher degree of suboptimality. Thus for the 3rd order plant model:

$$A = \begin{bmatrix} -0.33 & 0.31 & -0.14 \\ 0.31 & -0.53 & 0.07 \\ -0.13 & 0.07 & -0.04 \end{bmatrix} \quad \bar{B} = \begin{bmatrix} 1.61 \\ -0.12 \\ -3.31 \end{bmatrix} \quad B_1 = \begin{bmatrix} 1.80 \\ 1.20 \\ -0.80 \end{bmatrix} \quad B_2 = \begin{bmatrix} 1.40 \\ 0.20 \\ 1.60 \end{bmatrix}$$
$$c_1 = \begin{bmatrix} 0.80 \; 3.30 \; -3.20 \end{bmatrix} \quad c_2 = \begin{bmatrix} 2.60 \; 0.80 \; 1.20 \end{bmatrix},$$

with $U = 0.5$, $Y_2 = 2$, an approximate model realization involving only output map uncertainty can be constructed by identifying the means and variances of a pair of MA models. However, a stochastic MPC law for output map uncertainty designed using the approach of [12] (using 6th order MA models) gives an average closed-loop cost (over 100 initial conditions) of 114, whereas the average cost for (22) for the same set of initial conditions is 36.1.

References

[1] M.V. Kothare, V. Balakrishnan, and M. Morari. Robust constrained model predictive control using linear matrix inequalities. *Automatica*, 32:1361–1379, 1996.
[2] J.C. Allwright and G.C. Papvasiliou. Linear programming and robust model-predictive control using impulse-responses. *Systems & Control Letters*, 18:159–164, 1992.
[3] P.O.M. Scokaert and D.Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Trans. Autom. Control*, 43:1136–1142, 1998.
[4] A. Bemporad, F. Borrelli, and M. Morari. Min-max control of constrained uncertain discrete-time linear systems. *IEEE Trans. Automatic Control*, 48(9):1600–1606, 2003.
[5] D.W. Clarke, C. Mothadi, and P.S. Tuffs. Generalized predictive control. *Automatica*, 23:137–160, 1987.
[6] P. Li, M. Wendt, and G. Wozny. A probabilistically constrained model predictive controller. *Automatica*, 38(7):1171–1176, 2002.
[7] I. Batina. *Model predictive control for stochastic systems by randomized algorithms*. PhD thesis, Technische Universiteit Eindhoven, 2004.
[8] A.J. Felt. Stochastic linear model predictive control using nested decomposition. In *Proc. American Control Conf.*, pages 3602–3607, 2003.
[9] B. Kouvaritakis, M. Cannon, and P. Couchman. MPC as a tool for sustainable development integrated policy assessment. *IEEE Trans. Autom. Control*, 51(1):145–149, 2006.
[10] J.H. Lee and B.L. Cooley. Optimal feedback control strategies for state-space systems with stochastic parameters. *IEEE Trans. Autom. Control*, 43(10):1469–1475, 1998.
[11] D. 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 Control, Orlando*, 2001.
[12] P. Couchman, M. Cannon, and B. Kouvaritakis. Stochastic MPC with inequality stability constraints. *Automatica*, 2006. Submitted.
[13] P. Couchman, B. Kouvaritakis, and M. Cannon. LTV models in MPC for sustainable development. *Int. J. Control*, 79(1):63–73, 2006.
[14] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36:789–814, 2000.

[15] M. Cannon, B. Kouvaritakis, and V. Deshmukh. Enlargement of polytopic terminal region in NMPC by interpolation and partial invariance. *Automatica*, 40:311–317, 2004.
[16] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan. *Linear Matrix Inequalities in System and Control Theory*. SIAM, 1994.
[17] M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284:193–228, 1998.
[18] B.T. Polyak. *Introduction to optimization*. Opt. Software Inc., 1987.