Available online at www.sciencedirect.com
SCIENCE DIRECT®
Automatica 41 (2005) 595–604
automatica
www.elsevier.com/locate/automatica

Incorporating state estimation into model predictive control and its application to network traffic control☆

Jun Yan*, Robert R. Bitmead

Department of Mechanical & Aerospace Engineering, University of California, San Diego, 9500 Gilman Drive, La Jolla CA 92093-0411, USA

Received 10 July 2003; received in revised form 21 October 2004; accepted 4 November 2004

Abstract

Model predictive control (MPC) is of interest because it is one of the few control design methods which preserves standard design variables and yet handles constraints. MPC is normally posed as a full-state feedback control and is implemented in a certainty-equivalence fashion with best estimates of the states being used in place of the exact state. This paper focuses on exploring the inclusion of state estimates and their interaction with constraints. It does this by applying constrained MPC to a system with stochastic disturbances. The stochastic nature of the problem requires re-posing the constraints in a probabilistic form. Using a gaussian assumption, the original problem is approximated by a standard deterministically-constrained MPC problem for the conditional mean process of the state. The state estimates' conditional covariances appear in tightening the constraints. 'Closed-loop covariance' is introduced to reduce the infeasibility and the conservativeness caused by using long-horizon, open-loop prediction covariances. The resulting control law is applied to a telecommunications network traffic control problem as an example.

Keywords: Model predictive control; Network control; Constrained control

1. Introduction

Model predictive control (MPC) is an increasingly significant and popular control approach because of its use of a possibly nonlinear multivariable process model and its ability to handle constraints on inputs, states and outputs. It uses open-loop constrained optimization of finite-horizon control criteria in a receding horizon approach. A model is used to predict the future behavior of the system up to the horizon, starting from its current state, and a constrained optimization based on the prediction yields an optimal open-loop control sequence over the complete horizon. Only the first element in this sequence is applied to the plant. New measurements available at the next sample time permit the calculation of an updated initial state value, and the optimization is then re-solved. The introduction of each output measurement, via the mechanism of state update, results in the overall method yielding a closed-loop control.

The receding horizon approach behind MPC relies on state estimation even though most analyses of the stability, feasibility and performance of these schemes treat the controller as a full-state-feedback strategy (Mayne, Rawlings, Rao, & Scokaert, 2000). From MPC's early linear unconstrained variants such as generalized predictive control and its connection to LQG (Clarke, Mohtadi, & Tuffs, 1987; Bitmead, Gevers, & Wertz, 1990), the formulation of MPC has included state estimation either inherently or explicitly in the construction of the predictor using observer polynomials or via the Kalman filter. However, to our knowledge, there has been no satisfactory treatment of the inclusion of state

☆ This paper was not presented at any IFAC meeting. This paper was recommended for publication in revised form by Associate Editor I. Paschalidis under the direction of Editor I. Petersen. Research supported by USA National Science Foundation Grants ECS-0200449 and ECS-0225530.
* Corresponding author.
E-mail addresses: junyan@mae.ucsd.edu (J. Yan), rbitmead@mae.ucsd.edu (R.R. Bitmead).

0005-1098/$ - see front matter © 2005 Elsevier Ltd. All rights reserved.
doi:10.1016/j.automatica.2004.11.022

estimation into the formulation of the full constrained problem other than our own (Yan & Bitmead, 2002) and some very interesting recent work of van Hessen and Bosgra (2002), both of which we build on here. This is in spite of the industrial impetus for MPC being its application in process control, where the rejection of stochastic disturbances is the central control objective. The formulations cited above are certainty-equivalence controllers, in which the state estimate is used as a replacement for the exact state in the control law. The natural sequencing of the receding horizon control calculation would see the one-step-ahead state prediction used as the initial state in the computation of the MPC control sequence.

The core issue is that, when stochastic disturbances and/or model uncertainty come into the picture, the state estimates are at variance with the actual state. In this case, the conventional predictive controller cannot necessarily guarantee the fulfillment of state or output constraints in the real system, even though they might be satisfied for the predicted system. There are two sources of error involved: the error between the state estimate used and the actual state at the initial time of the optimization interval, and the future state errors introduced with the evolution of the system along its prediction horizon.

Our approach to accommodating the state estimate error into constrained MPC will involve three steps:

Step i: Replace the deterministic state and output constraints by probabilistic constraints.

Step ii: Use an approximate distribution for the state estimate error to convert these probabilistic constraints into deterministic constraints on the conditional mean of the state.

Step iii: Solve this standard deterministic constrained MPC problem for the future control sequence.

The first step weakens the hard constraints to a probabilistic form, at least for unbounded distributions. But this is necessary if, indeed, the disturbance process admits such signals. This modification has also been proposed by Li, Wendt, and Wozny (2000, 2002) and potentially leads to a rather intractable stochastic programming problem. The second step is the crux of our approach, where the state estimates and their covariances are introduced together to yield a noncertainty-equivalence MPC controller. The third step is computationally standard in MPC.

The novelties of our approach to incorporating state estimates into MPC are:

- introduction of the probabilistic constraints, building on some ideas of Yan and Bitmead (2002) and van Hessen and Bosgra (2002),
- inclusion of the state estimate's conditional mean and conditional covariance into the deterministic MPC formulation,
- development of the concepts of *closed-loop covariances*, in which the estimate covariances are moderated along the prediction horizon to reflect only minimal-delay values.

The net outcome of the approach is that the MPC with state estimates remains an MPC problem using the prediction of the state in place of the actual state. The *quality* of the state estimate, typically as measured by its covariance, is used to tighten the constraints in this modified deterministic MPC problem. This is a naturally appealing idea of compromising or de-tuning performance in order to augment robustness to stochastic or other disturbances.

Stochastic optimal control problems have been considered for some time (Kushner, 1971; Åström, 1970; Fleming & Rishel, 1975) and generically involve solution via dynamic programming in which the entire conditional distribution function evolves with time. Attempts to pose and solve probabilistically constrained optimal control problems have been made under the banner of MPC. The main task is to find a more simply soluble approximation of the original problem. For instance, Li et al. (2000, 2002) consider the input–output model of a linear system with normally distributed disturbances. An MPC strategy with probabilistic constraints on the output is posed and then is transformed into a deterministic nonlinear programming problem. Only open-loop prediction is used and the computational efficiency is also considered.

In van Hessen and Bosgra (2002), the authors used a special form of control composed of a nominal value and a linear feedback, in the probabilistically constrained MPC problem. The probability of a set of linear constraints on the state is replaced by an ellipsoidal approximation, yielding a conic optimization problem. A number of other recent articles, see e.g. Bemporad (1998) and Fukushima and Bitmead (2003, 2005) also treat MPC as a perturbation of a nominal stabilizing feedback and accommodate disturbances through tightening of the constraints. This is also an outcome of our approach, although here the focus will also be on maintaining an accessible MPC reformulation. Further, in the penultimate section, an example from telecommunications network traffic control will be developed, which illustrates the implementation and properties of the approach.

As will be shown in Section 3.1, under a certain gaussian assumption or approximation, the probabilistic constraints provide a convenient and natural point of entry for the estimate covariance. This is done by imposing stricter constraints reflecting the estimate quality. By including this extra information, the stochastic optimal control problem is converted to a soluble deterministic optimization problem in standard MPC form, in which the initial state is replaced by its one-step-ahead conditional-mean estimate and the constraints are modified to accommodate the state estimate's conditional covariance. This is equivalent to the original stochastic problem and is a noncertainty-equivalence controller, because of its dependence on the covariance.

The state prediction covariance normally is an increasing function of prediction horizon and so one might expect that constraint modification reflecting such a covariance would lead to a conservative control law or even infeasibility as the constraints tighten along the horizon, because the closed-loop property of the system is not taken into account. However, we know that future measurements will become available and, accordingly, in Section 4.2 we use the one-step-ahead covariance throughout the horizon to modify the constraints. We denominate this the *closed-loop covariance*. The resultant approximation problem can be easily solved by quadratic programming (QP) routines while the probabilistic feasibility is ensured.

2. MPC with probabilistic constraints

The archetypal deterministic MPC problem is as follows:

$$\begin{aligned}\min_u \quad & J(N, x_n, u_n^{n+N-1}) \\ & = x_{n+N}^{\mathrm{T}} P_N x_{n+N} + \sum_{i=0}^{N-1} (x_{n+i}^{\mathrm{T}} Q_i x_{n+i} \\ & \quad + u_{n+i}^{\mathrm{T}} R_i u_{n+i}) \\ \text{s.t.} \quad & x_{n+i+1} = f(x_{n+i}, u_{n+i}), \\ & x_{n+i} \in \mathbb{X}_i \ \ (i = 1, \ldots, N), \\ & u_{n+i} \in \mathbb{U}_i \ \ (i = 0, \ldots, N-1).\end{aligned}$$

The minimization commences at time n from initial state value x_n and yields the N-step solution sequence for time n, $u_n^{n+N-1} = \{u_n, u_{n+1}, \ldots, u_{n+N-1}\}$. Clearly output equations and constraints can be accommodated in this formulation, since the output depends explicitly on the state and input.

Now consider the following discrete-time stochastic system:

$$\begin{aligned} x_{n+1} &= f(x_n, u_n, \omega_n), \\ y_n &= g(x_n, u_n, v_n), \end{aligned}$$

where the first equation is the recursion of state variable x_n with process noise term ω_n and the second equation is that of the measurement y_n perturbed by noise v_n. Denote by $\mathscr{Y}_n$ the set of the collected input and output measurements $\{(u_i, y_i) : i = 1, \ldots, n\}$ and by $E_{\mathscr{Y}_n}(\cdot)$ and $P_{\mathscr{Y}_n}(\cdot)$ the conditional expectation and conditional distribution with respect to it.

The probabilistically constrained MPC problem we consider is

$$\begin{aligned}\text{MPC1}: \quad \min_u \quad & J(N, x_n, u_n^{n+N-1}) \\ & = E_{\mathscr{Y}_n}[x_{n+N}^{\mathrm{T}} P_N x_{n+N} \\ & \quad + \sum_{i=0}^{N-1} (x_{n+i}^{\mathrm{T}} Q_i x_{n+i} \\ & \quad + u_{n+i}^{\mathrm{T}} R_i u_{n+i})] \\ \text{s.t.} \quad & x_{n+i+1} = f(x_{n+i}, u_{n+i}, \omega_{n+i}), \\ & y_{n+i} = g(x_{n+i}, u_{n+i}, v_{n+i}), \\ & P_{\mathscr{Y}_n}(x_{n+i} \in \mathbb{X}_i) \geqslant p_i \ \ (i = 1, \ldots, N), \\ & u_{n+i} \in \mathbb{U}_i \quad (i = 0, \ldots, N-1).\end{aligned}$$

Remarks. (1) The criterion is the conditional expectation of a quadratic cost function of x and u conditioned on $\mathscr{Y}_n$. By minimizing this as part of a receding horizon strategy, we obtain a feedback control law.

(2) The constraints on state variables are posed in terms of probabilities, i.e. the constraints may be violated but only at a specific (presumably very low, $0 \leqslant (1 - p_i) \ll 1$) rate.

(3) If the disturbances belong to a compact set, then min–max or worst case methods might be applied. Two types of min–max MPC have been proposed: open- and closed loop. The open-loop strategy, as shown in Lee and Yu (1997), inevitably leads to an overly conservative control scheme. The closed-loop strategies (Lee & Yu, 1997; Scokaert & Mayne, 1998; Bemporad, 1998) can reduce the conservativeness but are computationally demanding.

(4) The constraints on the controls u_n^{n+N-1} in MPC1 are deterministic. Since they are free variables in the optimization, there is no need to put probabilistic constraints on them.

(5) This is a difficult constrained stochastic optimization problem (Fleming & Rishel, 1975). The solution would nominally involve the determination of the entire conditional distribution of the state and not just its first few moments.

(6) Our approach will be to take an alternative system based on a gaussian assumption or approximation of the state estimate error distribution. This will then permit a simple recasting of the probabilistic MPC problem. We shall require some fairly strong (but often acceptable) assumptions about the underlying system, notably that the future state covariance is independent of the current control. This is a property trivially satisfied for linear systems.

3. Incorporating state estimates into MPC

The task of this section is to seek a tractable deterministic approximation of the problem MPC1 of Section 2. For simplicity, we commence with a constrained linear system to develop the approach and later highlight the role played by linearity in our technique.

Consider a k-dimensional linear system with mutually independent gaussian initial condition $x_0 \sim N(\hat{x}_0, \Sigma_0)$, gaussian process noise $\omega_n \sim N(0, \Gamma_n)$ and gaussian measurement noise $v_n \sim N(0, \Lambda_n)$. Take the constraints in MPC1 as linear, giving

$$\begin{aligned}\text{MPC2}: \quad \min_u \quad & J(N, x_n, u_n^{n+N-1}) \\ & = E_{\mathscr{Y}_n}[x_{n+N}^{\mathrm{T}} P_N x_{n+N} \\ & \quad + \sum_{i=0}^{N-1} (x_{n+i}^{\mathrm{T}} Q_i x_{n+i} \\ & \quad + u_{n+i}^{\mathrm{T}} R_i u_{n+i})] \qquad (1) \\ \text{s.t.} \quad & x_{n+i+1} = A x_{n+i} + B u_{n+i} + G \omega_{n+i}, \qquad (2) \\ & y_{n+i} = C x_{n+i} + D v_{n+i}, \qquad (3) \\ & P_{\mathscr{Y}_n}(x_{n+i} \leqslant \beta_i) \geqslant p_i \ \ (i = 1, \ldots, N), \qquad (4) \\ & u_{n+i} \leqslant \mu_i \ \ (i = 0, \ldots, N-1). \qquad (5)\end{aligned}$$

3.1. Replacing probabilistic constraints by deterministic constraints

The problem at hand is that for system (2) at time n we want to solve the constrained stochastic optimal control problem MPC2. Ignoring for the moment the constraints, we have a linear system with gaussian disturbances ω_n and v_n. If the initial state distribution were gaussian and control were linear, then the state would also be gaussian, as would its conditional distribution function conditioned on $\mathscr{Y}_n$. We could then propagate the conditional distribution $P_{\mathscr{Y}_n}(x_{n+i})$ using the finite-dimensional Kalman filter:

$$\hat{x}_{n|n} = \hat{x}_{n|n-1} + L_n(y_n - C\hat{x}_{n|n-1}), \tag{6}$$

$$\hat{x}_{n+1|n} = A\hat{x}_{n|n} + Bu_n, \tag{7}$$

$$\Sigma_{n|n} = \Sigma_{n|n-1} - L_n C \Sigma_{n|n-1}, \tag{8}$$

$$\Sigma_{n+1|n} = A\Sigma_{n|n}A^{\mathrm{T}} + G\Gamma_n G^{\mathrm{T}}, \tag{9}$$

$$L_n = \Sigma_{n|n-1}C^{\mathrm{T}}(C\Sigma_{n|n-1}C^{\mathrm{T}} + \Lambda_n)^{-1}.$$

The Kalman filter and its predictor variants (see Anderson & Moore, 1979) compute the conditional means, $\hat{x}_{n+i|n}$, and conditional covariances $\Sigma_{n+i|n}$, which parametrize completely the gaussian conditional distribution functions.

We remark here, that since all the past control inputs $\{u_{n-i} : i = 0, 1, ..., n\}$ are known, assuming zero mean initial state estimate error, $\tilde{x}_{0|0} = x_0 - \hat{x}_{0|0}$, it follows that, by linearity, the filtered and predicted state estimates' errors, $\tilde{x}_{n+i|n} = x_{n+i} - \hat{x}_{n+i|n}$, are gaussian with means zero and covariances $\Sigma_{n+i|n}$, even though the feedback control might be nonlinear due to the constraints. [This is discernible from the Kalman filter error equations. We also note that the distribution of the state and of its estimate need not be gaussian for this to hold.] Further, the conditional state estimate error covariance is independent from the control. Thus, $\tilde{x}_{n+i|n} \sim N(0, \Sigma_{n+i|n})$. These conditional predicted means, prediction errors and covariances are computed from initial conditions $\hat{x}_{n|n}$, $\tilde{x}_{n|n}$ and $\Sigma_{n|n}$ using the open-loop predictor

$$\hat{x}_{n+i|n} = A\hat{x}_{n+i-1|n} + Bu_{n+i-1}, \tag{10}$$

$$\Sigma_{n+i|n} = A\Sigma_{n+i-1|n}A^{\mathrm{T}} + G\Gamma_n G^{\mathrm{T}} \ (i = 1, \ldots, N). \tag{11}$$

We use this to replace the probabilistic constraints of MPC2 by deterministic constraints on the conditional means. We draw attention to the additional property that these conditional means are optimized in the solution of the MPC control problem.

We first consider a scalar state system in modifying MPC2 into a deterministic form. This is as derived in Yan and Bitmead (2002). The multivariable case is a geometric extension using inscribed conic sets, see van Hessen and Bosgra (2002) and some later comments.

The gaussian distribution of the prediction errors makes it possible to modify constraints (4), $P_{\mathscr{Y}_n}(x_{n+i} \leqslant \beta_i) \geqslant p_i$, into a deterministic form. Since $\tilde{x}_{n+i|n} \sim N(0, \Sigma_{n+i|n})$, the derived variable $\xi_{n+i} = \Sigma_{n+i|n}^{-1/2}\tilde{x}_{n+i|n}$ has standard normal distribution $N(0, 1)$. Then,

$$P_{\mathscr{Y}_n}(x_{n+i} \leqslant \beta_i) \geqslant p_i \Leftrightarrow P_{\mathscr{Y}_n}(\xi_{n+i} \leqslant \check{\beta}_i) \geqslant p_i,$$

where, $\check{\beta}_i = \Sigma_{n+i|n}^{-1/2}(\beta_i - \hat{x}_{n+i|n})$. Denote by β_i^* the solution of $\Phi(\beta_i^*) = p_i$, where $\Phi(\cdot)$ is the standard normal distribution function. Then (4) can be recast as

$$P_{\mathscr{Y}_n}(x_{n+i} \leqslant \beta_i) \geqslant p_i \Leftrightarrow \hat{x}_{n+i|n} \leqslant \beta_i - \Sigma_{n+i|n}^{1/2}\beta_i^*. \tag{12}$$

The probabilistic constraints may thus be replaced by deterministic ones on the conditional mean process of the state.

The derivation has assumed that the state process is scalar, in order that the linear inequality constraint on x_{n+i} translates into a linear constraint on the $N(0, 1)$ variable ξ_{n+i}. For the vector state case, it is necessary to convert the set of linear constraints into an inscribed ellipsoidal constraint on the standard normal variables. This process is developed in more detail and in a similar context in van Hessen and Bosgra (2002). The end result is the conversion of the linear probabilistic constraints on x_{n+i} into tighter deterministic linear constraints on $\hat{x}_{n+i|n}$.

3.2. New deterministic MPC problem

Since the criterion function (1) is a conditional expectation, we are able to change it into a deterministic function of $\hat{x}_{n+i|n}$, u_{n+i} and $\Sigma_{n+i|n}$:

$$\begin{aligned}
J &= E_{\mathscr{Y}_n}[tr(x_{n+N}^{\mathrm{T}} P_N x_{n+N})] \\
&\quad + \sum_{i=0}^{N-1}\{E_{\mathscr{Y}_n}[tr(x_{n+i}^{\mathrm{T}} Q_i x_{n+i})] + u_{n+i}^{\mathrm{T}} R_i u_{n+i}\} \\
&= tr[E_{\mathscr{Y}_n}(x_{n+N}x_{n+N}^{\mathrm{T}})P_N] \\
&\quad + \sum_{i=0}^{N-1}\{tr[E_{\mathscr{Y}_n}(x_{n+i}x_{n+i}^{\mathrm{T}})Q_i] + u_{n+i}^{\mathrm{T}} R_i u_{n+i}\} \\
&= tr[(\Sigma_{n+N|n} + \hat{x}_{n+N|n}\hat{x}_{n+N|n}^{\mathrm{T}})P_N] \\
&\quad + \sum_{i=0}^{N-1}\{tr[(\Sigma_{n+i|n} + \hat{x}_{n+i|n}\hat{x}_{n+i|n}^{\mathrm{T}})Q_i] + u_{n+i}^{\mathrm{T}} R_i u_{n+i}\} \\
&= tr(\Sigma_{n+N|n}P_N) + \sum_{i=0}^{N-1} tr(\Sigma_{n+i|n}Q_i) \\
&\quad + \hat{x}_{n+N|n}^{\mathrm{T}} P_N \hat{x}_{n+N|n} \\
&\quad + \sum_{i=0}^{N-1}(\hat{x}_{n+i|n}^{\mathrm{T}} Q_i \hat{x}_{n+i|n} + u_{n+i}^{\mathrm{T}} R_i u_{n+i}).
\end{aligned} \tag{13}$$

From (11) the state prediction covariance $\Sigma_{n+i|n}$ will not be changed by the controls. Therefore, minimizing (13) is

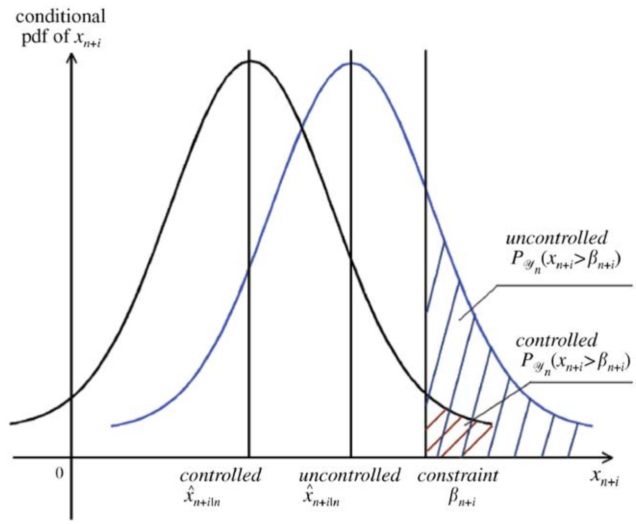

Fig. 1. Controlling the conditional mean process to achieve the probabilistic constraint.

equivalent to minimizing

$$J'(N, \hat{x}_{n|n}, u_n^{n+N-1}) = \hat{x}^{\mathrm{T}}_{n+N|n} P_N \hat{x}_{n+N|n} + \sum_{i=0}^{N-1} (\hat{x}^{\mathrm{T}}_{n+i|n} Q_i \hat{x}_{n+i|n} + u^{\mathrm{T}}_{n+i} R_i u_{n+i}).$$

Hence we have a new receding horizon control problem to be solved at each time instant.

$$\begin{aligned} \text{MPC3}: \quad \min_u \quad & J'(N, \hat{x}_{n|n}, u_n^{n+N-1}) \\ & = \hat{x}^{\mathrm{T}}_{n+N|n} P_N \hat{x}_{n+N|n} + \sum_{i=0}^{N-1} (\hat{x}^{\mathrm{T}}_{n+i|n} Q_i \hat{x}_{n+i|n} + u^{\mathrm{T}}_{n+i} R_i u_{n+i}) \\ \text{s.t.} \quad & \hat{x}_{n+i|n} \leqslant \beta_i - \Sigma^{1/2}_{n+i|n} \beta_i^* \ (i = 1, \ldots, N), \\ & u_{n+i} \leqslant \mu_i \ (i = 0, \ldots, N-1). \end{aligned} \tag{14}$$

The overall control strategy is the receding horizon implementation of this. As the state estimate covariance is used to modify the constraints, the control u is not certainty equivalence. Fig. 1 depicts two conditional probability density functions of the scalar state x_{n+i}. Their means are two differing values of $\hat{x}_{n+i|n}$ and they share a common variance. The probabilistic constraint represents the requirement that the shaded area under the conditional density curve to the right of β_{n+i} in Fig. 1 should be small enough. The new MPC task is to control the conditional mean values of state x to a proper level so that the probabilistic constraint can be achieved. The transformations and calculations earlier are designed to achieve this via (12).

Despite the depiction in Fig. 1, it is not true that x is gaussian. Because with the presence of the constraints, the control u is a nonlinear function of x or $\hat{x}$, as shown in Maciejowski (2002). Hence, the closed-loop system has nonlinear dynamics of state propagation and the gaussian property cannot be preserved in such a nonlinear process. However, the prediction error remains gaussian, which is the underpinning of our method.

As mentioned earlier, MPC3 is derived for a scalar system. In this case, we may state the following theorem.

Theorem 1. *For a linear scalar system in the form of* (2) *with measurement* (3), *if the initial estimation error and the random disturbances are independent zero-mean gaussian variables, then the minimizing solution of* MPC3 *solves* MPC2.

The transformation from MPC2 to MPC3 could be explained as that we tighten the constraints posed on the model according to the estimate quality (measured by the covariance of the estimation error) to ensure the satisfaction by those on the plant. The covariance of the open-loop prediction that evolves as in (11) has the effect of modifying the constraints. Typically, this covariance is increasing with i and in (12) this leads to successively more stringent restrictions being applied on the control. For large horizon N, feasibility could be lost. Even if there exists a feasible solution for the problem, that solution might be very conservative. The next section will be devoted to proposing a remedy for this difficulty using closed-loop covariances instead of the open-loop ones.

If the system is multi-dimensional, then (4) represents a joint probability of keeping the state in a polyhedron in state space. To evaluate the probability of vector constraints requires the calculation of multiple integrals of multi-dimensional normal density function, this is computationally demanding. It is convenient to blend the estimate covariance into the constraints. We propose to do the same ellipsoidal approximation as derived in van Hessen and Bosgra (2002). Replacing the probability of the polyhedron on the left-hand side of (4) by that of the maximal inscribed ellipsoid, provides an overbound of (4) which is suitable to apply with the gaussian assumption. However, the solution turns out to be conservative. Another option is to pose the constraints on the states in ellipsoidal form, if that makes sense, in (4) from the beginning.

To derive our deterministic MPC problem, we have chosen a linear system framework. Since MPC is attractive because of its ability to handle nonlinearities, it behoves us to comment on the dependence on this property. Its importance lies in establishing the quality of the gaussian approximation to the state, which permits the conversion of a stochastic problem to a deterministic one. Clearly, if an overbounding gaussian variable can be found for the state then this might be admissible in constraint management. The linearity also has been used to develop adequate descriptions of the distributions, via the Kalman filter or other nonlinear estimators.

This suggests that it might be possible to consider nonlinear systems operating with the extended Kalman filter. For our method, the central property of linear systems is that the state estimate covariance does not depend on the control signal chosen—this helps separate the constraint specification from the control solution. In a nonlinear context this would require other tricks to bound.

4. Closed-loop covariances

4.1. Covariance properties along the horizon

The modified constraints of (14) in MPC3 are increasingly stringent. The number β_i^* is positive (since p_i should represent a large probability) and fixed by the original constraint. Computation of $\Sigma_{n+i|n}$ proceeds according to the time-update portion of the Kalman filter (9). For the linear, stationary system earlier, this is given by (11)

$$\Sigma_{n+i+1|n} = A\Sigma_{n+i|n}A^{\mathrm{T}} + G\Gamma_{n+i}G^{\mathrm{T}},$$

and is initialized by $\Sigma_{n+1|n}$, the one-step-ahead prediction covariance. In the asymptotically stationary case, we have

$$\Sigma_{n+i|n} \geqslant \Sigma_{n+i-1|n} \geqslant \cdots \geqslant \Sigma_{n+1|n}. \quad (15)$$

The net result of inequalities such as (15) is that incorporation of the state estimate covariance into MPC, as developed in Section 3, implies possibly increasingly demanding modification of the constraints on the states as one moves along the horizon. Eventually, unless Σ_∞ is strongly bounded, infeasibility of the MPC problem will result through the growth of $\Sigma_{n+i|n}$ for large values of i. This is problematic because it would appear to militate against long horizons in MPC, which normally are associated with improved dynamical properties.

4.2. Closed-loop covariance values

Our solution to the conundrum of ever-increasingly tighter constraints along the horizon of MPC is to fix all the state estimate covariances at their one-step-ahead, or minimal-control-delay values. That is, the constraint re-formulation from (14),

$$\hat{x}_{n+i|n} \leqslant \beta_i - \Sigma_{n+i|n}^{1/2}\beta_i^*,$$

is replaced by its one-step-ahead variant

$$\hat{x}_{n+i|n} \leqslant \beta_i - \Sigma_{n+1|n}^{1/2}\beta_i^*.$$

We denominate this substitution as using the *closed-loop covariance*.

The reasoning behind the closed-loop covariance lies in the recognition that a receding horizon control strategy is being used. Thus, we know that before the next step of the MPC is taken a new measurement will become available. Thus, before the constraint involving $\Sigma_{n+2|n}$ might need to be active, a new measurement at time $n+1$ will have been available and MPC3 will be reposed and resolved. The control u_{n+1} that feeds into the plant will be computed subject to the constraint with bound $(\beta_2 - \Sigma_{n+2|n+1}^{1/2}\beta_2^*)$, which would be equal to $(\beta_2 - \Sigma_{n+1|n}^{1/2}\beta_2^*)$ in the stationary case. From Section 4.1, this quantity is no smaller than $(\beta_2 - \Sigma_{n+2|n}^{1/2}\beta_n^*)$. This argument may be repeated up to time $n+N$.

Therefore, we reach our final version of MPC strategy

$$\begin{aligned}
\text{MPC4}: \quad \min_u \quad & J'(N, \hat{x}_{n|n}, u_n^{n+N-1}) \\
&= \hat{x}_{n+N|n}^{\mathrm{T}} P_N \hat{x}_{n+N|n} \\
&\quad + \sum_{i=0}^{N-1} (\hat{x}_{n+i|n}^{\mathrm{T}} Q_i \hat{x}_{n+i|n} \\
&\quad + u_{n+i}^{\mathrm{T}} R_i u_{n+i}) \\
\text{s.t.} \quad & \hat{x}_{n+i|n} \leqslant \beta_i - \Sigma_{n+1|n}^{1/2}\beta_i^* \\
& (i = 1, \ldots, N), \\
& u_{n+i} \leqslant \mu_i \ (i = 0, \ldots, N-1).
\end{aligned}$$

Remarks.

- First we make two simple observations:
  - (i) The feasible region defined in MPC4 is typically bigger than that in MPC3.
  - (ii) The criteria in MPC4 and MPC3 are the same function. From (i), we can conclude that the optimized criterion functions satisfy $J_{\mathrm{MPC4}}^{\prime\mathrm{opt}} \leqslant J_{\mathrm{MPC3}}^{\prime\mathrm{opt}}$.
- In a sense, the use of the closed-loop covariance is dual to the receding-horizon control idea. While a sequence of N control values is computed, only the first is applied. The state update then converts this open-loop control into closed-loop and issues such as asymptotic stability are addressable. Similarly, while the prediction occurs over N time steps, only the one-step-ahead prediction covariance information plays a role in modifying the constraints. This again reflects the disguised closed-loop nature of the control and leads to the improved treatment of feasibility.
- Asymptotic stability of the system under receding-horizon MPC control and in the absence of disturbances or state estimation errors can be guaranteed provided the sequence of finite-horizon MPC problems includes a terminal state constraint applied on each finite horizon (Mayne et al., 2000). Similar methods also can be used to guarantee feasibility of the sequence of constrained problems.
- For linear systems with gaussian noises, guaranteed stability and feasibility are not possible, because of the unboundedness of the distribution. This will cause the disturbance to drive the state arbitrarily far from the constraint set. However, for systems with bounded noises, properly posed MPC4 can provide a mechanism to over-bound the probability of large excursion and hence ensure the satisfaction of the constraints.

5. A network traffic control example

In this section, the control of available bit rate (ABR) traffic in an asynchronous transfer mode (ATM) telecommunications network is considered as an example of applying our MPC method, MPC4. This is a multi-source, single-buffer congestion control problem.

5.1. ABR congestion control

The ABR service plays a central role in regulating telecommunications network traffic. The current standards for ATM traffic management are built on the foundation of a rate-based (rate matching) flow control scheme; see details in Imer, Compans, Başar, and Srikant (2001). The goal of ABR service congestion control is to provide fairness among all links with a minimal cell loss ratio and maximal utilization of network sources. The control challenge is to regulate ABR source rates to utilize maximally the available capacity as it varies while respecting the requirement not to overflow buffer queues too frequently.

We apply MPC to the ABR control problem. The downstream ABR is stochastically varying and the limited buffer queue size provides an obvious state constraint. The dynamics consist of the queue's integral action, stochastic variation of the ABR and action delays in the response of controlled data sources to commands from the node. This problem has been studied in a similar formulation by a number of authors (Mascolo, 1999; Altman, Başar, & Srikant, 1999; Imer & Başar, 1999; Imer et al., 2001). Our modeling framework and criterion are basically the same as in Altman et al. (1999), explicitly taking delay into account, and treating the available bit rate service as an autoregressive (AR) process driven by white noise. The action delays and the AR process are incorporated into the states by augmenting the system dimension. Consequently, the system model is posed in the form of (2) and we introduce an explicit probabilistic constraint on the queue length.

5.2. MPC problem formulation

The ABR control problem is depicted in Fig. 2. Main data streams are shown as solid lines, while network control data is shown as dotted lines. All signals are shown at the time of arrival or departure from the node. At time n at the node, it receives the current ABR, μ_n, from downstream and determines the source rate allocations, $v_{m,n}$, for each source $m = 1, \ldots, M$. The arriving data rates from the sources are $r_{m,n} = v_{m,n-d_m}$, where the action delay d_m for each source is due to the round trip transmission time from the node. The queue length is denoted q_n and evolves according to

$$q_{n+1} = sat_{[0,1]}\left(q_n + \sum_{m=1}^{M} r_{m,n} - \mu_n\right),$$

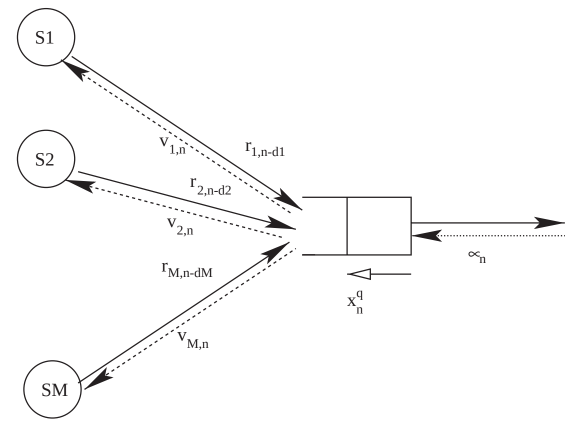

Fig. 2. ABR control problem.

where $sat_{[0,1]}(\cdot)$ is the saturation function and $[0, 1]$ corresponds to the range of the queue length from empty to full. It is presumed that the ABR is the sum with saturation of a known constant nominal service rate, μ, plus a pth order autoregressive process, ξ_n.

$$\mu_n = sat_{[0,1]}(\mu + \xi_n),$$
$$\xi_n = \sum_{i=1}^{p} \alpha_i \xi_{n-i} + \phi_{n-1},$$

where $\{\phi_n\}$, an i.i.d., $N(0, \sigma^2)$ distributed sequence, and α_i, $i = 1, \ldots, p$ are known coefficients.

The effective source rate $r_{m,n}$ is the response of data source m to a commanded rate or control action from the node d_m samples earlier. This action delay accounts for transmission time from the node to the source and return. Denoting the node control action for source m as $v_{m,n}$ we have, $v_{m,n-d_m} = r_{m,n}$.

Define the centered variables $x_n^{\rm q} = q_n - \bar{Q}$, $u_{m,n} = v_{m,n} - (1/M)\mu$, where $\bar{Q}$ is the nominal target mean queue utilization and μ is the mean ABR value.

To apply MPC4 to this nonlinear system, we stick with the linear part of the system as the model running for the MPC controller. First introduce $M-1$ new states

$$x^{\rm c}_{n+1} = I_M\, x^{\rm c}_n + \begin{pmatrix} 1 & -1 & 0 & \ldots & 0 \\ 0 & 1 & -1 & & 0 \\ \vdots & & & \vdots & \\ 0 & 0 & \ldots & 1 & -1 \end{pmatrix} \begin{pmatrix} u_{1,n} \\ u_{2,n} \\ \vdots \\ u_{M,n} \end{pmatrix}, \qquad (16)$$

which integrate the differences between the control inputs and will be included in the criterion function to achieve fairness among the sources. Then this formulation of the ABR control problem may be written in state-space form as (2)

$$x_{n+1} = Ax_n + Bu_n + G\omega_n,$$

where the state $x_n = (x_n^{\mathrm{dT}} \ x_n^{\mathrm{aT}} \ x_n^{\mathrm{cT}})^{\mathrm{T}}$ with x_n^{d} the state of the transmission delays, x_n^{a} the state of the autoregressive ABR model. Vector u_n is the input vector appearing in (16). The detailed formulation is similar to that developed in Yan and Bitmead (2002) except for the queue dynamics: $x_{n+1}^{\mathrm{q}} = x_n^{\mathrm{q}} + x_n^{\mathrm{d},1} - sat_{[-\mu,1-\mu]}(\xi_n)$. Note that, in this model for MPC, the saturation of the queue length is omitted, since the constraints should cope with the upper bound and the lower violation can be avoided by setting the nominal queue length $\bar{Q}$ and the bound of constraints away from it. For simplicity in development of the state estimator, we assume that at each sample time the real value of ξ is available.

The objective function to be minimized is given by

$$J = \frac{1}{N} E_{\mathscr{Y}_n} \left\{ \sum_{i=1}^{N} \left[x_{n+i}^{\mathrm{qT}} x_{n+i}^{\mathrm{q}} + \delta \sum_{j=1}^{M} \left(u_{j,n+i} - \frac{1}{M} sat_{[-\mu,1-\mu]}(\xi_{n+i+d_j}) \right)^2 + \gamma x_{n+i}^{\mathrm{cT}} x_{n+i}^{\mathrm{c}} \right] \right\}. \quad (17)$$

The first term in this criterion is the performance objective of keeping the queue close to an operating point. The second term tries to ensure that all available capacity is used. The third term manages fairness by making the average values of assigned rates the same.

Next, a set of probabilistic constraints, related to the upper bound on the queue length, is introduced:

$$P_{\mathscr{Y}_n}(q_{n+i} \leqslant Q_{\max}) = P_{\mathscr{Y}_n}(x_{n+i}^{\mathrm{q}} \leqslant Q_{\max} - \bar{Q}) \geqslant p_i \quad (i = d_{\min}+1, \ldots, N). \quad (18)$$

This constraint says that to some small extent losing and re-transmitting a certain proportion of data is tolerable—p_i should be chosen close to 1. Integer $d_{\min}$ denotes the smallest delay among the sources. Due to the time delay, at time n states $x_n^{\mathrm{q}} \cdots x_{n+d_{\min}}^{\mathrm{q}}$ are not controllable, therefore, they should be constraint-free.

Now apply the method developed earlier by modeling $\tilde{x}_{n+i}^{\mathrm{q}}$ and $\{\xi_n\}$ as gaussian. Constraint (18) may be replaced by

$$\hat{x}_{n+i|n}^{\mathrm{q}} \leqslant Q_{\max} - \bar{Q} - \beta_i^* \Sigma_{n+d_{\min}+1|n}^{\mathrm{q}^{1/2}},$$

where $\Phi(\beta_i^*) = p_i$, and $\Phi(\cdot)$ is the standard normal distribution function. Note, that we have used the closed-loop covariance $\Sigma_{n+d_{\min}+1|n}^{\mathrm{q}}$ for all values of i. As $\tilde{x}$ involves the saturation of the AR process, its distribution is not normal but a truncated version of the normal distribution function. However, the constraints still work, because $\Sigma_{n+d_{\min}+1|n}^{\mathrm{q}}$ is an overbound of the real covariance. Finally, the MPC problem becomes

$$\text{ABR– MPC}: \quad \min_u \quad J = \frac{1}{N} E_{\mathscr{Y}_n} \left\{ \sum_{i=1}^{N} \left[x_{n+i}^{\mathrm{qT}} x_{n+i}^{\mathrm{q}} + \delta \sum_{j=1}^{M} \left(u_{j,n+i} - \frac{1}{M} sat_{[-\mu,1-\mu]}(\xi_{n+i+d_j}) \right)^2 + \gamma x_{n+i}^{\mathrm{cT}} x_{n+i}^{\mathrm{c}} \right] \right\} \quad (19)$$

$$\text{s.t.} \quad \hat{x}_{n+i|n}^{\mathrm{q}} \leqslant Q_{\max} - \bar{Q} - \beta_i^* \Sigma_{n+d_{\min}+1|n}^{\mathrm{q}^{1/2}} \quad (i = d_{\min}+1, \ldots, N). \quad (20)$$

And the control sent into the system should be $sat_{[-1/M,1-1/M]}(u_n^j)$, since any real source rate can neither be negative nor exceed the maximum available service rate.

For the purpose of comparison, we also run the simulation for the stabilizing controller suggested in Imer and Başar (1999) which is an infinite horizon LQG controller with the criterion function being the limit of (17) without saturation in the second term and the fairness term $\gamma x^{\mathrm{cT}} x^{\mathrm{c}}$. As proved in Imer and Başar (1999), this approach can guarantee a stabilizing control law. The simulation result shows that fairness among the sources cannot be obtained by the second term in the criterion function. Another LQG controller is also simulated for the same system with the same criterion function (17) as in ABR–MPC and the controls are subject to the saturation. This receding horizon approach is a fair strategy. A comparison of these three controller will be presented next.

5.3. Simulation results

5.3.1. Results of ABR–MPC

The simulation is of a single congested node accessed by three sources with delays $d_1 = 2, d_2 = 4, d_3 = 6$. The AR process is assumed to be second order with parameters $\alpha_1 = 1.88$, $\alpha_2 = -0.8836$. Coefficients δ and γ are both set to 0.0001. The finite horizon should be taken bigger than the biggest action delay, so $N = 10$. The nominal queue length, $\bar{Q}$ in the criterion function, is set to 2. This means the constraints are always active. The probabilistic constraint is $P(x_{n+i}^{\mathrm{q}} \leqslant -0.2) \geqslant 0.9505$, with corresponding $\beta^* = 1.65$. The total simulation time is 10,000 s. Fig. 3 is part of the plot of the available service rate μ_n and the queue q_n. In this simulation result, there are about 488 steps in 10,000 overflowed, which satisfies the expected portion 5%. The means of the source rates $r_{1,n}, r_{2,n}, r_{3,n}$ are 0.1634, 0.1634, 0.1634, respectively, indicating that fairness has been achieved. The variances of these source rates are 0.0223, 0.0198, 0.0189, which is decreasing with increasing round-trip delay from the node. This is because the lesser delayed source rate is more effective in regulating the queue length, and therefore is used more frequently to manage the overflow.

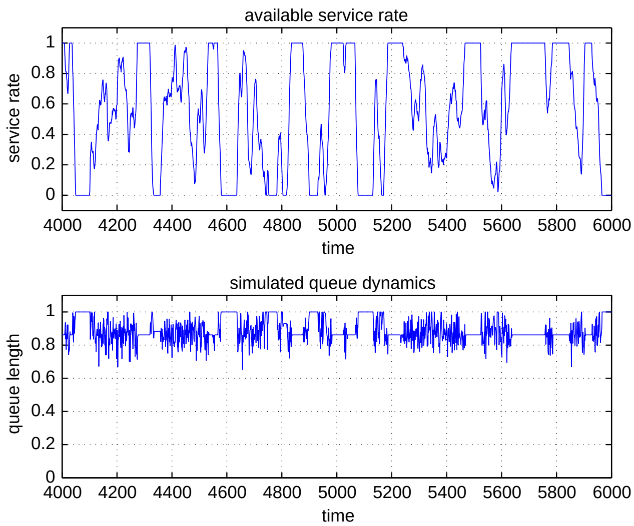

Fig. 3. Available service rate and the queue length for the ABR–MPC control.

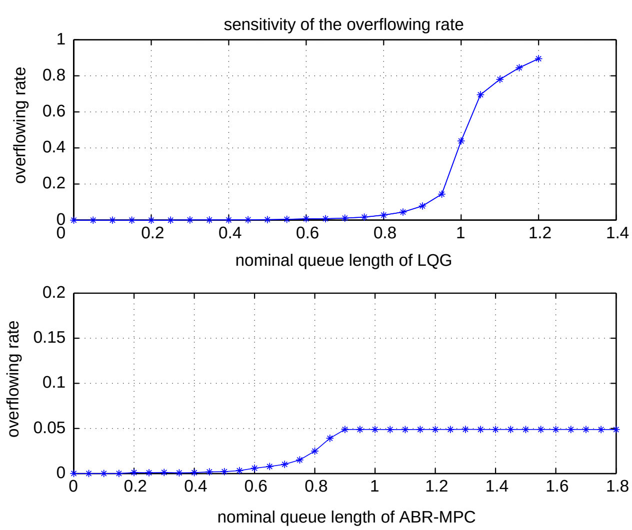

Fig. 4. Sensitivity to nominal queue length (note the difference in scales).

5.3.2. Comparison of existing approaches and ABR–MPC

The infinite horizon LQG controller proposed in Imer and Başar (1999) and the receding horizon LQG controller are simulated with the same noise set applied to ABR–MPC for 10,000 samples. The nominal queue length $\bar{Q}$ is the only tuning parameter for these two strategies and is set to 0.65 and 0.85, respectively, to achieve 5% overflowing rate, the primary performance requirement. The means of the source rates given by the infinite horizon LQG are 0.1720,0.1584 and 0.1562. Therefore this is not a fair strategy. The receding horizon LQG is a fair strategy with the means of the sources rates 0.1634, the same as ABR–MPC, hence these two strategies almost have the same total data throughput.

Though the two controllers, receding horizon LQG and ABR–MPC, have the same performance, LQG is sensitive to the nominal queue length $\bar{Q}$ while ABR–MPC is robust to it. To show this, we simulated these two controllers for different values of $\bar{Q}$, 10,000 samples for each. The result is shown in Fig. 4, the plot of the overflowing rate for different nominal queue values. It can be seen that it is difficult for LQG to achieve the exact 5% overflowing rate via tuning $\bar{Q}$. For ABR–MPC, the overflowing rate can be guaranteed by setting the nominal queue length great enough, i.e. make the constraints active which were designed corresponding to the given overflowing rate. Once some parameters of the system are changed (e.g. the covariance of the white noise), LQG might require re-tuning $\bar{Q}$ which cannot be done explicitly; while ABR–MPC would require re-calculating and change of the corresponding terms in the constraints directly.

6. Conclusion

This paper has presented an approach to the treatment of state estimates in constrained MPC. The key ideas are:

- that the constraints should be converted to a probabilistic form,
- that the probabilistic constraint on the states should be replaced by a deterministic constraint on the controlled conditional mean state estimate and that this modified constraint should involve tighter stipulations reflecting the estimate covariance,
- that the closed-loop covariance should be used for the modification of all constraints along the horizon.

The resulting deterministic MPC problem is standard in its form but uses the state estimate as its starting state and introduces covariance information via the constraints. This is not a certainty-equivalence controller.

The ideas are demonstrated on a recent telecommunications network control problem of interest and found to display an ease of use and effectiveness of application in this almost linear, constrained problem. A key feature leading to this facility was the property of linear systems that the state estimates' covariances are independent of the control signal.

Further work remains to be tackled in extending the approach to a wide range of more obviously nonlinear problems. Equally, there would be considerable advantage to developing a method to overbound simply probabilities of constraint violation with normal approximations. This is the focus of current work.

References

Altman, E., Başar, T., & Srikant, R. (1999). Congestion control as a stochastic control problem with action delays. *Automatica*, *35*, 1937–1950.

Anderson, B. D., & Moore, J. B. (1979). *Optimal filtering*. Englewood Cliffs, NJ: Prentice-Hall.

Åström, K. J. (1970). *Introduction to stochastic control*. New York: Academic Press.

Bemporad, A. (1998). Reducing conservativeness in predictive control of constrained systems with disturbances. In *Proceedings of the 37th IEEE conference on decision and control*, Tampa, FL, USA.

Bitmead, R., Gevers, M., & Wertz, V. (1990). *Adaptive optimal control—the thinking man's GPC*. New York: Prentice-Hall.

Clarke, D., Mohtadi, C., & Tuffs, P. (1987). Generalized predictive control. part 1: the basic algorithms. *Automatica, 23*, 137–148.

Fleming, W. H., & Rishel, R. W. (1975). *Deterministic and stochastic optimal control*. Berlin: Springer.

Fukushima, H., & Bitmead, R. (2003). Robust constrained model predictive control using closed-loop prediction. In *Proceedings of the 42nd IEEE conference on decision and control*, Maui, HI, USA.

Fukushima, H., & Bitmead, R. (2005). Robust constrained model predictive control using comparison model. *Automatica, 41*, 97–106.

Imer, O., & Başar, T. (1999). Optimal solution to a team problem with informaiton delays: an application in flow control for communication networks. In *Proceedings of the 38th conference on decision and control*, Phoenix, AZ, USA.

Imer, O., Compans, S., Başar, T., & Srikant, R. (2001). Available bit rate congestion control in atm networks. *IEEE Control Systems Magazine*, *135*, 38–56.

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

Lee, J., & Yu, Z. (1997). Worst-case formulation of model predictive control for systems with bounded parameters. *Automatica*, *33*, 763–781.

Li, P., Wendt, M., & Wozny, G. (2000). Robust model predictive control under chance constraints. *Computers and Chemical Engineering*, *24*, 829–834.

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

Maciejowski, J. (2002). *Predictive control with constraints*. Englewood Cliffs, NJ: Prentice-Hall.

Mascolo, S. (1999). Congestion control in high-speed communication networks using the smith principle. *Automatica*, *35*, 1921–1935.

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

Scokaert, P., & Mayne, D. (1998). Min–max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic control*, *43*, 1136–1142.

van Hessen, D., & Bosgra, O. (2002). A conic reformulation of model predictive control including bounded and stochastic disturbances under state and input constraints. In *Proceedings of the 41st IEEE conference on decision and control*, Las Vegas, NV, USA.

Yan, J., & Bitmead, R. R. (2002). Model predictive control and state estimation: a network example. In *Proceedings of the 15th IFAC*, Barcelona, Spain.

Jun Yan is pursuing his Ph.D. degree with the Dynamic Systems and Control group, Department of Mechanical and Aerospace Engineering, University of California, San Diego since 2000. He received his M.S. degree from Department of Mechanical and Aerospace Engineering, University of California, San Diego in 2002 and his B.S. degree from the Department of Automation, Tsinghua University, Beijing, China in 2000.

Robert Bitmead holds the Cymer Corporation Chair at the Department of Mechanical and Aerospace Engineering, University of California, San Diego. He received his Ph.D. from the University of Newcastle in 1980. He was with the Australian National University Department of Systems Engineering from 1982–1999. He is a Fellow of the Australian Academy of Technological Sciences and Engineering and a Fellow of the IEEE. He was member of the IFAC Council from 1996–2002 and a member of IFAC Technical Board from 1990–1996. He hails originally from Sydney and was General Chair of the IEEE Conference on Decision and Control there from December 12–15, 2000. He is Editor of Automatica for Adaptive and Intelligent Control.