

Application of Interior-Point Methods to Model Predictive Control¹

C. V. RAO,² S. J. WRIGHT,³ AND J. B. RAWLINGS⁴

Communicated by D. Q. Mayne

Abstract. We present a structured interior-point method for the efficient solution of the optimal control problem in model predictive control. The cost of this approach is linear in the horizon length, compared with cubic growth for a naive approach. We use a discrete-time Riccati recursion to solve the linear equations efficiently at each iteration of the interior-point method, and show that this recursion is numerically stable. We demonstrate the effectiveness of the approach by applying it to three process control problems.

Key Words. Model predictive control, interior-point methods, Riccati equation.

1. Introduction

Model predictive control (MPC) is an optimal control-based strategy that uses a plant model to predict the effect of an input profile on the evolving state of the plant. At each step of MPC, an optimal control problem with Bolza objectives is solved and its optimal input profile is implemented until another plant measurement becomes available. The updated plant

¹This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under Contract W-31-109-Eng-38, and a grant from Aspen Technology. We acknowledge the support of the industrial members of the Texas-Wisconsin Modeling and Control Consortium. We are grateful to three referees of the original version of this paper, whose insightful comments improved the paper considerably.

²Research Assistant, Department of Chemical Engineering, University of Wisconsin, Madison, Wisconsin.

³Computer Scientist, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Illinois.

⁴Professor, Department of Chemical Engineering, University of Wisconsin, Madison, Wisconsin.

information is used to formulate and solve a new optimal control problem, thereby providing feedback from the plant to the model, and the process is repeated. This strategy yields a receding horizon control formulation.

The MPC methodology is appealing to the practitioner because input and state constraints can be explicitly accounted for in the controller. A practical disadvantage is its computational cost, which has tended to limit MPC applications to linear processes with relatively slow dynamics. For such problems, the optimal control problem to be solved at each stage of MPC is a convex quadratic program. While robust and efficient software exists for the solution of unstructured convex quadratic programs, significant improvements often can be made by exploiting the structure of the MPC subproblem.

When input and state constraints are not present, MPC with an infinite horizon is simply the well-known linear-quadratic regulator problem. Even when constraints are present, the infinite-horizon MPC problem reduces generally to a linear-quadratic regulator after a certain number of stages (see Refs. 1–3) and therefore can be recast as a finite-dimensional quadratic program. Since this quadratic program can be large, with many stages, it is important that algorithms be efficient for problems with long horizons.

Unconstrained discrete-time linear-quadratic optimal control problems can be solved by using a discrete-time Riccati equation. The computational cost of this algorithm is linear in the horizon length N . A different formulation, obtained by eliminating the state variables, results in an unconstrained quadratic function whose Hessian is dense, with dimensions that grow linearly in N . The cost of minimizing this quadratic function is cubic in N , making it uncompetitive with the Riccati approach in general. There is a third option, an optimization formulation in which the states are retained explicitly as unknowns in the optimization and the model equation is retained as a constraint. The optimality conditions for this formulation reveal that the adjoint variables are simply the Lagrange multipliers for the model equation and that the problem can be solved by factoring a matrix whose dimension again grows linearly with N . In this formulation, the matrix is banded, with a bandwidth independent of N , so the cost of the factorization is linear rather than cubic in N . The discrete-time Riccati equation can be interpreted as a block factorization scheme applied to this matrix.

Traditionally, the discrete-time Riccati equation is obtained by using dynamic programming to solve the unconstrained linear optimal control problem. The essential idea in dynamic programming is to work stage-by-stage through the problem in reverse order, starting with the final stage N . The optimization problem reduces to a simpler problem at each stage; see Berksekas (Ref. 4) for further details. Block factorization, like dynamic programming, exploits the multistaged nature of the optimization problem.

The key difference is that the block factorization approach tackles the problem explicitly, whereas dynamic programming tackles the problem semi-implicitly by using the Bellman principle of optimality. The explicit treatment allows greater flexibility, since the block factorization approach retains its inherent structure even when inequality constraints are added to the formulation.

When constraints are present, the scheme for unconstrained problems must be embedded in an algorithmic framework that determines which of the inequalities are active and which are inactive at the optimum. At each iteration of the outer algorithm, the main computational operation is the solution of a set of linear equations whose structure is very like that encountered in the unconstrained problem. Hence, the cost of performing each iteration of the outer algorithm is linear in the number of stages N . This observation has been made by numerous authors, in the context of outer algorithms based on both active-set and interior-point methods. Glad and Johnson (Ref. 5) and Arnold et al. (Ref. 6) demonstrate that the factorization of a structured Lagrangian in an optimal control problem with a Bolza objective for an active set framework yields a Riccati recursion. Wright (Refs. 7–8), Steinback (Ref. 9), and Lim et al. (Ref. 10) investigate the Bolza control problem in an interior-point framework.

In this paper, we present an MPC algorithm based on an interior-point method, in which a block factorization is used at each iteration to obtain the search direction for the interior-point method. Our work differs from earlier contributions in that the formulation of the optimal control problem is tailored to the MPC application, the interior-point algorithm is based on Mehrotra's algorithm (Ref. 11), whose practical efficiency on general linear and quadratic programming problems is well documented, and the linear system at each interior-point iteration is solved efficiently by a Riccati recursion. We compare our approach with the alternative of using the model equation to eliminate the states, yielding a dense quadratic program in the input variables alone, and present results obtained for three large industrial problems.

We use order notation in the following (standard) way: If a matrix, vector, or scalar quantity M is a function of another matrix, vector, or scalar quantity E , we write

$$M = O(\|E\|), \quad \text{if there is a constant } \beta \text{ such that } \|M\| \leq \beta \|E\|, \\ \text{for all } \|E\| \text{ sufficiently small.}$$

We write

$$M = \Theta(\|E\|), \quad \text{if there is a constant } \beta \text{ such that} \\ \|E\|/\beta \leq \|M\| \leq \beta \|E\|.$$

We say that a matrix is positive diagonal if it is diagonal with positive diagonal elements. The term nonnegative diagonal is defined correspondingly. We use SPD as an abbreviation for symmetric positive definite and SPSD as an abbreviation for symmetric positive semidefinite.

2. Model Predictive Control

2.1. Infinite-Horizon Problem. The fundamental formulation of the linear model predictive controller is the following infinite-dimensional convex quadratic program:

$$\min_{u, x} \Phi(u, x) = (1/2) \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k + \Delta u_k^T S \Delta u_k), \quad (1)$$

subject to the following constraints:

$$x_0 = \hat{x}_j, \quad x_{k+1} = A x_k + B u_k, \quad (2a)$$

$$D u_k \leq d, \quad G \Delta u_k \leq g, \quad H x_k \leq h, \quad (2b)$$

where

$$x_k \in \mathbb{R}^n, \quad u_k \in \mathbb{R}^m, \quad \Delta u_k = u_k - u_{k-1}.$$

The vector \hat{x}_j represents the current estimate of the state at discrete time j , whereas x_k represents the state at k sampling steps along the future prediction horizon and u_k represents the input at this same time. We assume that Q and S are SPSD matrices and that R is SPD.

By a suitable adjustment of the origin, the formulation (1)–(2) can also account for target tracking and disturbance rejection (Ref. 12). If there is a feasible point for the constraints (2), the infinite-horizon regulator formulation is stabilizing whenever (A, B) is stabilizable and $(A, Q^{1/2})$ is detectable (Ref. 13).

For unstable state transition matrices, (1)–(2) is ill-conditioned because the infinite-horizon formulation can potentially yield unbounded solutions. To improve the conditioning of the optimization, we parameterize the input as

$$u_k = L x_k + r_k,$$

where L is a linear stabilizing feedback gain for (A, B) ; see Refs. 14–15. The system model becomes

$$x_{k+1} = (A + BL)x_k + B r_k, \quad (3)$$

where r_k is the new manipulated input. By initially specifying a stabilizing (but potentially infeasible) trajectory, we can improve the numerical conditioning of the optimization by excluding unstable solutions.

By expanding Δu_k , we transform (1)–(2) into the following more tractable form:

$$\min_{u, x} \Phi(u, x) = (1/2) \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k), \quad (4)$$

subject to

$$x_0 = \hat{x}_j, \quad x_{k+1} = A x_k + B u_k, \quad (5a)$$

$$D u_k - G x_k \leq d, \quad H x_k \leq h. \quad (5b)$$

The original formulation (1)–(2) can be recovered from (4)–(5) by making the following substitutions into the second formulation:

$$\begin{aligned} \hat{x}_j &\leftarrow \begin{bmatrix} \hat{x}_j \\ u_{j-1} \end{bmatrix}, & x_k &\leftarrow \begin{bmatrix} x_k \\ u_{k-1} \end{bmatrix}, & u_k &\leftarrow r_k, \\ A &\leftarrow \begin{bmatrix} A + BL & 0 \\ L & 0 \end{bmatrix}, & B &\leftarrow \begin{bmatrix} B \\ I \end{bmatrix}, & Q &\leftarrow \begin{bmatrix} Q + L^T(R + S)L & -L^T S \\ -SL & S \end{bmatrix}, \\ M &\leftarrow \begin{bmatrix} L^T(R + S) \\ -S \end{bmatrix}, & R &\leftarrow R + S, & D &\leftarrow \begin{bmatrix} D \\ G \end{bmatrix}, \\ G &\leftarrow \begin{bmatrix} -DL & 0 \\ -L & G \end{bmatrix}, & d &\leftarrow \begin{bmatrix} d \\ g \end{bmatrix}, & H &\leftarrow [H, 0]. \end{aligned}$$

In the remainder of this section, we address two issues. The first is the replacement of (4)–(5) by an equivalent or similar finite-horizon problem, a step necessary for the practical computation of the solution. The second issue is the replacement of the constraints $Hx_k \leq h$ by so-called soft constraints. Instead of enforcing these conditions strictly, we add terms to the objective that penalize violations of these conditions. This technique is a more appropriate way of dealing with certain constraints from an engineering point of view.

2.2. Receding Horizon Regulator Formulation. The key step in reducing (4)–(5) to a finite-horizon problem is the use of a linear control law to determine u_k after a certain time horizon, that is,

$$u_k = K x_k, \quad \text{for all } k \geq N. \quad (6)$$

With this added constraint, the states x_k , $k > N$, and the inputs u_k , $k \geq N$, are completely determined by x_N , the state at the end of the prediction horizon.

Two techniques can be used to determine the law (6). The first, due to Rawlings and Muske (Ref. 16), sets $K=0$ uniformly in (6) and produces an approximate solution to (4)–(5). With this substitution, we have

$$(1/2) \sum_{k=N}^{\infty} (x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k) = (1/2) \sum_{k=N}^{\infty} x_k^T Q x_k. \quad (7)$$

If A is stable, this sum is equal to $(1/2)x_N^T \bar{Q} x_N$, where \bar{Q} is the solution of the matrix Lyapunov equation

$$\bar{Q} - A^T \bar{Q} A = Q. \quad (8)$$

If A is unstable, the sum (7) may be infinite, so we impose a stabilizing constraint to derive any useful information from the solution of the model problem. The Schur decomposition of A can be used to construct a basis for the stable subspace of A . If this decomposition is

$$A = UTU^T = [U_s, U_u] \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix} \begin{bmatrix} U_s^T \\ U_u^T \end{bmatrix},$$

where the eigenvalues of T_{11} are inside the unit circle whereas those of T_{22} are contained on or outside the unit circle, then the (orthogonal) columns of U_s span the stable subspace of A and the (orthogonal) columns of U_u span the orthogonal complement of the stable subspace of A . We add the endpoint constraint

$$F x_N = 0, \quad \text{where } F = U_u^T, \quad (9)$$

to ensure that the unstable modes have vanished by stage N (Ref. 17). Since the input u_k is zero for all $k \geq N$, the unstable modes also remain at zero at all subsequent stages. The evolution of the stable modes on the infinite horizon can be accounted for by solving the following Lyapunov equation for \bar{Q} :

$$\bar{Q} - A_s^T \bar{Q} A_s = Q, \quad \text{where } A_s = U_s T_{11} U_s^T,$$

and replacing the infinite sum with $(1/2)x_N^T \bar{Q} x_N$, as above.

In the second formulation, discussed in Refs. 1–3 and 14, the input after stage N is parameterized with the classical linear quadratic gain K obtained from the solution of the steady-state Riccati equation. This matrix, used in conjunction with the control law (6), is the solution to the unconstrained version of the problem, in which the inequality constraints

(5b) do not appear. By using (6), the infinite tail of the sum in (4) can be written as

$$\begin{aligned} & (1/2) \sum_{k=N}^{\infty} (x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k) \\ & = (1/2) \sum_{k=N}^{\infty} (x_k^T Q x_k + x_k^T K^T R K x_k + 2x_k^T M K x_k). \end{aligned}$$

This infinite summation can be replaced by the single term $(1/2)x_N^T \bar{Q} x_N$, where \bar{Q} is the solution of the following discrete-time algebraic Riccati equation:

$$\bar{Q} = Q + A^T \bar{Q} A - (A^T \bar{Q} B + M)(R + B^T \bar{Q} B)^{-1} (B^T \bar{Q} A + M^T). \quad (10)$$

In both formulations, the feedback law (6) is valid only if the constraints (5) are satisfied at all stages, including the stages $k \geq N$. Hence, we would like to implement this law only after we reach a state x_N such that the solution generated by the control law (6) and the model equation in (5a) at stages $k \geq N$ satisfies the inequalities (5b) at all such stages. We define a set \mathcal{X} of states for which this property holds as follows:

$$\mathcal{X} = \{x: H(A - BK)^l x \leq h, (DK - G)(A - BK)^l x \leq d, \text{ for all } l \geq 0\},$$

where K is the optimal unconstrained linear control law obtained from the following equation:

$$K = -(R + B^T \bar{Q} B)^{-1} (B^T \bar{Q} A + M^T). \quad (11)$$

If N is chosen so that $x_N \in \mathcal{X}$, then the following finite-horizon problem is equivalent⁵ to (4)–(5):

$$\min_{u, x} \Phi(u, x) = (1/2) \sum_{k=0}^{N-1} (x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k) + (1/2)x_N^T \bar{Q} x_N, \quad (12)$$

subject to

$$x_0 = \hat{x}_j, \quad x_{k+1} = Ax_k + Bu_k, \quad k = 0, 1, 2, \dots, \quad (13a)$$

$$u_k = Kx_k, \quad k \geq N, \quad (13b)$$

$$Du_k - Gx_k \leq d, \quad Hx_k \leq h, \quad k = 0, 1, 2, \dots, N-1, \quad (13c)$$

where \bar{Q} is defined in (10).

⁵The finite-horizon problem is also a valid approximation to (4)–(5) for $K=0$ when the endpoint constraint $Fx_N=0$ is added to (13) and \bar{Q} is defined by (8).

If the components of h and d are strictly positive and if the unconstrained model is stabilizable, we can show that \mathcal{X} contains 0 in its interior. Under these circumstances, there is an index N_∞ such that

$$x_k \notin \mathcal{X}, \quad k < N_\infty, \quad x_k \in \mathcal{X}, \quad k \geq N_\infty. \quad (14)$$

Since N_∞ is difficult to obtain in practice and since \mathcal{X} is characterized by a finite number of conditions, we can solve the problem (12)–(13) for some fixed value of N and then check that the states and inputs at stages $k \geq N$ continue to satisfy the inequality constraints at subsequent stages. If not, we increase N and repeat the process.

A variety of methods guarantee that the constraints are satisfied on the infinite horizon by checking a finite number of stages $k \geq N$. Sockaert and Rawlings (Ref. 2) propose constructing an open ball B_x contained within the set \mathcal{X} , thereby allowing termination of the search when $x_k \in B_x$ for $k \geq N$. The approximation for \mathcal{X} tends to be conservative, since the algorithm is motivated by norm-bounding arguments. A more practical method, given by Gilbert and Tan (Ref. 18), is to construct explicitly the set \mathcal{X} . The constructive representation of \mathcal{X} provides a priori an upper bound l on the number of feasible stages $k \in [N, N+l]$ necessary to guarantee that all of the subsequent stages $k \geq l+N$ are feasible. The drawback of this approach is that the algorithm for constructing the maximal sets is not guaranteed to converge for unbounded feasible regions, since \mathcal{X} may be unbounded. For a compact, convex set of states, an alternative approach that circumvents having to check for constraint violations is given by Chmielewski and Manousiouthakis (Ref. 1). By examining the extremal points on the feasible region, they calculate a conservative upper bound on the N required to guarantee that the solution is feasible on the infinite horizon.

We have assumed to this point that there exists a feasible solution with respect to the input and endpoint constraints for the optimal control calculation. In the presence of the side constraints (2b), it is no longer true that the constrained regulator stabilizes all possible states even when the stabilizability assumption is satisfied. When stabilization is not possible, the problem (4)–(5) is an infeasible optimization problem. In actual operation, an infeasible solution would signal a process exception condition.

For the Rawlings–Muske formulation, enforcement of the endpoint constraint (9) often results in an infeasible optimization problem. Feasibility can be recovered often by increasing the horizon length N ; but when the initial state is not stabilizable, the feasible region will continue to be empty for all N . The existence of a feasible N can be checked easily by solving the following linear program:

$$\min_{u, x, r} e^T r, \quad (15)$$

where e is the vector whose entries are all 1, subject to the constraints

$$x_{k+1} = Ax_k + Bu_k, \quad k = 0, 1, 2, \dots, N-1, \quad (16a)$$

$$Du_k - Gx_k \leq d, \quad Hx_k \leq h, \quad k = 0, 1, 2, \dots, N-1, \quad (16b)$$

$$r - Fx_N \geq 0, \quad r + Fx_N \geq 0. \quad (16c)$$

A positive solution to the linear program indicates that a feasible solution does not exist and the horizon length N must be increased. If the feasibility check fails for some user-supplied upper bound on the horizon length, then the current state is not constrained stabilizable for the specified regulator.

2.3. Feasibility and Soft Constraints. In the formulation of the MPC problem, some state constraints are imposed by physical limitations such as valve saturation. Other constraints are less important; for instance, they may represent desired ranges of operation for the plant. In some situations, no set of inputs and states for the MPC problem may satisfy all of these constraints. Rather than having the algorithm declare infeasibility and return without a result, we prefer a solution that enforces some constraints strictly (hard constraints), while relaxing others and replacing them with penalties on their violation (soft constraints).

Scokaert and Rawlings (Ref. 13) replace the soft constraints with penalty terms in the objective that are a combination of l_1 -norms and squared l_2 -norms of the constraint violations. Assuming for simplicity that all state constraints $Hx_k \leq h$ in (13) are softened in this way, we obtain the following modification to the objective (12):

$$\begin{aligned} \min_{u, x, \epsilon} \Phi(u, x, \epsilon) = & (1/2) \sum_{k=0}^{N-1} (x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k + \epsilon_k^T Z \epsilon_k) \\ & + z^T \epsilon_k + (1/2)x_N^T \bar{Q} x_N, \end{aligned} \quad (17)$$

where the constraint violations ϵ_k are defined by the following formulas, which replace $Hx_k \leq h$:

$$Hx_k - \epsilon_k \leq h, \quad \epsilon_k \geq 0, \quad (18)$$

and the elements of the vector z are nonnegative, while Z is an SPSD matrix. It is known that, when the weighting z on the l_1 -terms is sufficiently large [see, for example, Section 12.3 in Fletcher (Ref. 19)], and when the original problem (12)–(13) has a nonempty feasible region, the local minimizers of problem (12)–(13) modified by (17)–(18) defined above correspond to local solutions of the unmodified problem (12)–(13). Under these conditions, the formulation (17) together with the constraints (18) is referred to as an

exact penalty formulation of the original objective (12) with the original constraints $Hx_k \leq h$. This formulation has the advantage that it can still yield a solution when the original problem (12)–(13) is infeasible.

Prior to actually solving the problem, we cannot know how large the elements of z must be chosen to make the exact penalty property hold. The threshold value depends on the optimal multipliers for the original problem (12)–(13). A conservative state-dependent upper bound for these multipliers can be obtained by exploiting the Lipschitz continuity of the quadratic program (Ref. 20). In practice, the exact penalty is not critical, since by definition soft constraints need not be satisfied exactly. Reasonable controller performance can be achieved often by setting $z=0$ and choosing Z to be a positive diagonal matrix. In fact, the inclusion of the l_2 -term $\epsilon_k^T Z \epsilon_k$ is not needed at all for the exact penalty property to hold, but is included here to provide a little more flexibility in the modeling.

In the remainder of the paper, we work with a general form of the MPC problem, which contains all the features discussed in this section: finite horizon, endpoint constraints, and soft constraints. This general form is

$$\begin{aligned} \min_{u, x, \epsilon} \Phi(u, x, \epsilon) = & \sum_{k=0}^{N-1} (1/2)(x_k^T Q x_k + u_k^T R u_k + 2x_k^T M u_k + \epsilon_k^T Z \epsilon_k) \\ & + z^T \epsilon_k + (1/2)x_N^T \bar{Q}_N x_N, \end{aligned} \quad (19)$$

subject to

$$x_0 = \hat{x}_j, \quad (20a)$$

$$x_{k+1} = Ax_k + Bu_k, \quad k = 0, 1, \dots, N-1, \quad (20b)$$

$$Du_k - Gx_k \leq d, \quad k = 0, 1, \dots, N-1, \quad (20c)$$

$$Hx_k - \epsilon_k \leq h, \quad k = 1, 2, \dots, N, \quad (20d)$$

$$\epsilon_k \geq 0, \quad k = 1, 2, \dots, N, \quad (20e)$$

$$Fx_N = 0. \quad (20f)$$

Note that (20a) is fixed and that (20f) is required only when we choose the parameterization $K=0$. We assume throughout that the matrices in (19) satisfy the properties

$$R \text{ is PSD, } \quad Z \text{ is SPSD, } \quad \begin{bmatrix} Q & M \\ M^T & R \end{bmatrix} \text{ is SPSD.} \quad (21)$$

Note that the last property holds for the matrices considered in Section 2.1, since in making the substitutions to obtain the form (4) we obtain

$$\begin{aligned} \begin{bmatrix} Q & M \\ M^T & R \end{bmatrix} &\leftarrow \begin{bmatrix} Q+L^T(R+S)L & -L^T S & L^T(R+S) \\ -SL & S & -S \\ (R+S)L & -S & (R+S) \end{bmatrix} \\ &= \begin{bmatrix} Q & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} + \begin{bmatrix} L^T \\ -I \\ I \end{bmatrix} S[L, -I, I] + \begin{bmatrix} L^T \\ 0 \\ I \end{bmatrix} R[L, 0, I], \end{aligned}$$

which is a sum of SPSD matrices and is therefore itself SPSD.

3. Interior-Point Method

In this section, we describe our interior-point-based approach for solving the MPC problem (19)–(20). We start with a general description of the interior-point method of choice for linear and convex quadratic programming: Mehrotra’s predictor-corrector algorithm. The remaining sections describe the specialization of this approach to MPC, including the use of the Riccati approach to solve the linear subproblem, handling of endpoint constraints, and hot starting.

3.1. Mehrotra’s Predictor-Corrector Algorithm. Active set methods have proved to be efficient for solving quadratic programs with general constraints. The interior-point approach has proved to be an attractive alternative when the problems are large and convex. In addition, this approach has the advantage that the system of linear equations to be solved at each iterate has the same dimension and structure throughout the algorithm, making it possible to exploit any structure inherent in the problem. The most widely used interior-point algorithms do not require a feasible starting point to be specified. In fact, they usually generate infeasible iterates, attaining feasibility only in the limit. From a theoretical viewpoint, interior-point methods exhibit polynomial complexity, in contrast to the exponential complexity of active-set approaches.

In this section, we sketch an interior-point method for general convex quadratic programming problems and discuss its application to the specific problem (19). A more complete description is given by Wright (Ref. 21).

Consider the following convex quadratic program:

$$\min_w \Phi(w) = (1/2)w^T Qw + c^T w, \quad \text{s.t. } Fw = g, \quad Cw \leq d, \quad (22)$$

where Q is an SPSD matrix. The Karush–Kuhn–Tucker (KKT) conditions for optimality are that there exist vectors π^* and λ^* such that the following conditions are satisfied for $(w, \pi, \lambda) = (w^*, \pi^*, \lambda^*)$:

$$\begin{aligned} Qw + F^T \pi + C^T \lambda + c &= 0, \\ -Fw + g &= 0, \\ -Cw + d &\geq 0, \\ \lambda &\geq 0, \\ \lambda_j(-Cw + d)_j &= 0, \quad j = 1, 2, \dots, m, \end{aligned}$$

where m is the number of rows in the matrix C . Because the objective function is convex, the KKT conditions are both necessary and sufficient for optimality. By introducing a vector t of slacks for the constraint $Cw \leq d$, we can rewrite these conditions in a slightly more convenient form,

$$\mathcal{F}(w, \pi, \lambda, t) = \begin{bmatrix} Qw + F^T \pi + C^T \lambda + c \\ -Fw + g \\ -Cw - t + d \\ T\Lambda e \end{bmatrix} = 0, \quad (23a)$$

$$(\lambda, t) \geq 0, \quad (23b)$$

where T and Λ are diagonal matrices defined by

$$T = \text{diag}(t_1, t_2, \dots, t_m), \quad \Lambda = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_m),$$

and where

$$e = (1, 1, \dots, 1)^T.$$

Primal-dual interior-point methods generate iterates $(w^i, \pi^i, \lambda^i, t^i)$, $i = 1, 2, \dots$, with $(\lambda^i, t^i) > 0$ that approach feasibility with respect to the conditions (23a) as $i \rightarrow \infty$. The search directions are Newton-like directions for the equality conditions in (23a). Dropping the superscript and denoting the current iterate by (w, π, λ, t) , we can write the general linear system to be solved for the search direction as

$$\begin{bmatrix} Q & F^T & C^T & & \\ -F & & & & \\ -C & & & -I & \\ & & T & \Lambda & \end{bmatrix} \begin{bmatrix} \Delta w \\ \Delta \pi \\ \Delta \lambda \\ \Delta t \end{bmatrix} = \begin{bmatrix} r_Q \\ r_F \\ r_C \\ r_t \end{bmatrix}. \quad (24)$$

Note that the coefficient matrix is the Jacobian of the nonlinear equations (23a). Different primal-dual methods are obtained from different choices of the right-hand side vector (r_Q, r_F, r_C, r_t) . The duality gap μ , defined by

$$\mu = \lambda^T t / m, \tag{25}$$

is typically used as a measure of optimality of the current point (w, π, λ, t) . In principle, primal-dual interior-point methods ensure that the norm of the function \mathcal{F} defined by (23a) remains bounded by a constant multiple of μ at each iterate, thus ensuring that μ is also a measure of infeasibility of the current point. However, the latter condition is rarely checked in practical algorithms.

We use a variant of the Mehrotra predictor-corrector algorithm (Ref. 11) to solve (22). This algorithm has proved to be the most effective approach for general linear programs and is similarly effective for convex quadratic programming. The first part of the Mehrotra search direction, the predictor step or affine-scaling step, is simply a pure Newton step for the system (23a), obtained by solving (24) with the following right-hand side:

$$\begin{bmatrix} r_Q \\ r_F \\ r_C \\ r_t \end{bmatrix} = -\mathcal{F}(w, \pi, \lambda, t) = - \begin{bmatrix} Qw + F^T \pi + C^T \lambda + c \\ -Fw + g \\ -Cw - t + d \\ T\Lambda e \end{bmatrix}. \tag{26}$$

We denote the corresponding solution of (24) by $(\Delta w_{\text{aff}}, \Delta \pi_{\text{aff}}, \Delta \lambda_{\text{aff}}, \Delta t_{\text{aff}})$. The second part of the search direction, the centering-corrector direction $(\Delta w_{\text{cc}}, \Delta \pi_{\text{cc}}, \Delta \lambda_{\text{cc}}, \Delta t_{\text{cc}})$, is calculated by choosing the centering parameter $\sigma \in [0, 1)$ as outlined below and solving the system (24) with the following right-hand side:

$$\begin{bmatrix} r_Q \\ r_F \\ r_C \\ r_t \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ -\Delta T_{\text{aff}} \Delta \Lambda_{\text{aff}} e + \sigma \mu e \end{bmatrix}, \tag{27}$$

where ΔT_{aff} and $\Delta \Lambda_{\text{aff}}$ are the diagonal matrices constructed from the elements of Δt_{aff} and $\Delta \lambda_{\text{aff}}$, respectively.

The following heuristic for choosing the value of σ has proved to be highly effective. We first compute the maximum steplength α_{aff} that can be

taken along the affine-scaling direction as follows:

$$\alpha_{\text{aff}} = \arg \max\{\alpha \in [0, 1] \mid (\lambda, t) + \alpha(\Delta\lambda_{\text{aff}}, \Delta t_{\text{aff}}) \geq 0\}.$$

The duality gap μ_{aff} attained from this full step to the boundary is

$$\mu_{\text{aff}} = (\lambda + \alpha\Delta\lambda_{\text{aff}})^T (t + \alpha\Delta t_{\text{aff}}) / m.$$

Finally, we set

$$\sigma = (\mu_{\text{aff}} / \mu)^3.$$

The search direction is obtained by adding the predictor and centering-corrector direction as follows:

$$\begin{aligned} (\Delta w, \Delta \pi, \Delta \lambda, \Delta t) &= (\Delta w_{\text{aff}}, \Delta \pi_{\text{aff}}, \Delta \lambda_{\text{aff}}, \Delta t_{\text{aff}}) \\ &\quad + (\Delta w_{\text{cc}}, \Delta \pi_{\text{cc}}, \Delta \lambda_{\text{cc}}, \Delta t_{\text{cc}}). \end{aligned} \quad (28)$$

Note that the coefficient matrix in (24) is the same for both the predictor and centering-corrector systems, so just one factorization of this matrix is required at each iteration. Apart from this factorization, the main computational operations at each iteration include two back-substitutions for two different right-hand sides and a number of matrix-vector operations.

The distance along the direction (28) is defined in terms of the maximum step α_{max} that can be taken without violating the condition (23b),

$$\alpha_{\text{max}} = \arg \max\{\alpha \in [0, 1] \mid (\lambda, t) + \alpha(\Delta\lambda, \Delta t) \geq 0\}.$$

The actual steplength α is chosen to be

$$\alpha \leftarrow \gamma \alpha_{\text{max}}, \quad (29)$$

where γ is a parameter in the range (0, 1) chosen to ensure that the pairwise products $\lambda_i t_i$ do not become too unbalanced. The value of γ is typically close to 1; it has proved effective in practice to allow it to approach 1 as the algorithm gets closer and closer to the solution. See Mehrotra (Ref. 11) for the details of a heuristic for choosing γ .

The algorithm does not require the initial point to be feasible, and checks can be added to detect problems for which no feasible points exist. In our case, feasibility of the MPC problem obtained from the Rawlings and Muske formulation with unstable plants can be determined a priori by solving the linear program (15)–(16).

Finally, we note that block elimination can be applied to the system (24) to obtain reduced systems with more convenient structures. By eliminating Δt , we obtain the following system:

$$\begin{bmatrix} Q & F^T & C^T \\ -F & & \\ -C & & \Lambda^{-1}T \end{bmatrix} \begin{bmatrix} \Delta w \\ \Delta \pi \\ \Delta \lambda \end{bmatrix} = \begin{bmatrix} r_Q \\ r_F \\ r_C + \Lambda^{-1}r_t \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} \hat{r}_Q \\ \hat{r}_F \\ \hat{r}_C \end{bmatrix}. \quad (30)$$

Since $\Lambda^{-1}T$ is a positive diagonal matrix, we can eliminate easily $\Delta \lambda$ as well to obtain

$$\begin{bmatrix} Q + C^T \Lambda T^{-1} C & F^T \\ -F & 0 \end{bmatrix} \begin{bmatrix} \Delta w \\ \Delta \pi \end{bmatrix} = \begin{bmatrix} r_Q - C^T T^{-1} (\Lambda r_C + r_t) \\ r_F \end{bmatrix}. \quad (31)$$

As we see in the next section, these eliminations can be applied to our particular problem to put the system in a form in which we can apply the Riccati block-elimination technique of Sections 3.3 and 3.4.

We conclude with a note on the sizes of the elements in t and λ and their effect on the elements of the matrices in (30) and (31). In path-following interior-point methods that adhere rigorously to the theory, iterates are confined to a region in which the pairwise products $t_i \lambda_i$ are not too different from each other in size. A bound of the form

$$t_i \lambda_i \geq \gamma \mu \quad (32)$$

is usually enforced, where μ is the average value of $t_i \lambda_i$ [see (25)] and $\gamma \in (0, 1)$ is constant, typically $\gamma = 10^{-4}$. When the primal-dual solution set for (22) is bounded, we have further that

$$t_i \leq \beta, \quad \lambda_i \leq \beta, \quad i = 1, 2, \dots, m, \quad (33)$$

for some constant bound $\beta > 0$. It follows immediately from (32) and (33) that

$$(\gamma/\beta^2)\mu \leq t_i/\lambda_i \leq (\beta^2/\gamma)\mu^{-1}, \quad (34a)$$

$$(\gamma/\beta^2)\mu \leq \lambda_i/t_i \leq (\beta^2/\gamma)\mu^{-1}. \quad (34b)$$

Hence, the diagonal elements of the matrices $T^{-1}\Lambda$ and $\Lambda^{-1}T$ lie in the range $[\Theta(\mu), \Theta(\mu^{-1})]$.

Although bounds of the form (32) are not enforced explicitly in most implementations of the Mehrotra algorithm, computational experience shows that they are almost always satisfied in practice. Hence, it is reasonable to assume, as we do in the analysis of numerical stability below, that the estimates (34) are satisfied by iterates of our algorithm.

3.2. Efficient MPC Formulation. The optimal control problem (19)–(20) has been viewed traditionally as a problem in which just the inputs are variables, while the states are eliminated by direct substitution using the transition equation (20b); see, for example, Muske and Rawlings (Ref. 12). We refer to this formulation hereafter as the standard method. Unfortunately, the constraint and Hessian matrices in the reduced problem resulting from this procedure are generally dense, so the computational cost of solving the problem is proportional to N^3 . Efficient commercial solvers for dense quadratic programs [such as QPSOL (Ref. 22)] can then be applied to the reduced problem.

Unless the number of stages N is small, the $O(N^3)$ cost of the standard method is unacceptable because the unconstrained version of (19) is known to be solvable in $O(N)$ time by using a Riccati equation or dynamic programming. We are led to ask whether there is an algorithm for the constrained problem (19)–(20) that preserves the $O(N)$ behavior. In fact, the interior-point algorithm of the preceding section almost attains this goal, since it can be applied to the problem (19)–(20) at a cost of $O(N)$ operations per iteration. The rows and columns of the reduced linear systems (30) and (31) can be rearranged to make these matrices banded, with dimension proportional to N and bandwidth independent of N . Since the number of iterations required by the interior-point algorithm depends only weakly on N in practice, the total computational cost of this approach is only slightly higher than $O(N)$. In both the active set and interior-point approaches, the dependence of solution time on other parameters, such as the number of inputs, number of states, and number of side constraints, is cubic.

Wright (Refs. 7–8) describes a scheme in which these banded matrices are explicitly formed and factored with a general banded factorization routine. In the next section, we show that the linear system to be solved at each interior-point iteration can be reduced to a form identical to the unconstrained version of (19)–(20), that is, a form in which the side constraints (20c), (20d) are absent. Hence, a Riccati recursion similar to the technique used for the unconstrained problem can be used to solve this linear system. Even though such a scheme places restrictions on the use of pivoting for numerical stability, we show by a simple argument that numerical stability can be expected.

Suppose that the interior-point algorithm of Section 3.1 is applied to the problem (19)–(20). We use λ_k , ζ_k , η_k to denote the Lagrange multipliers for the constraints (20c), (20d), (20e), respectively. We rearrange the linear system (30) to be solved at each iteration of the interior-point method by interleaving the variables and equations according to stage index. That is, the primal and dual variables for stage 0 are listed before those for stage 1, and so on. For this ordering, the rows of the system (30) that correspond

and initial state) is a final point condition:

$$\min_{u, x} \Phi(u, x) = (1/2)u_0^T R_0 u_0 + \sum_{k=1}^{N-1} (1/2)(x_k^T Q_k x_k + u_k^T R_k u_k + 2x_k^T M_k u_k) + (1/2)x_N^T Q_N x_N, \tag{41}$$

subject to

$$x_0 = \hat{x}_j, \tag{42a}$$

$$x_{k+1} = Ax_k + Bu_k, \quad k=0, 1, \dots, N-1, \tag{42b}$$

$$Fx_N = 0. \tag{42c}$$

Note that (42a) is fixed.

The problem (41)–(42) is convex if the matrices R_0 , Q_N , and

$$\begin{bmatrix} Q_k & M_k \\ M_k^T & R_k \end{bmatrix}, \quad k=1, 2, \dots, N-1, \tag{43}$$

are all SPSD. The following brief discussion shows that this property holds. First, we show that

$$(\Sigma_k^H)^{-1} - (\Sigma_k^H Z_k \Sigma_k^H)^{-1} \text{ is positive diagonal for all } k=1, 2, \dots, N. \tag{44}$$

By using the definition of Z_k above, together with the diagonality of Z and Σ_k^ϵ , we have that

$$\begin{aligned} (\Sigma_k^H)^{-1} - (\Sigma_k^H Z_k \Sigma_k^H)^{-1} &= (\Sigma_k^H)^{-1} [I - (\Sigma_k^H Z_k)^{-1}] \\ &= (\Sigma_k^H)^{-1} [I - (\Sigma_k^H Z + \Sigma_k^H (\Sigma_k^\epsilon)^{-1} + I)^{-1}]. \end{aligned}$$

Since Z , Σ_k^H , Σ_k^ϵ are all positive diagonal matrices, the final expression above is a product of two positive diagonal matrices, and therefore is itself positive diagonal. Hence, property (44) holds. From (39a), note that Q_N is an SPSD modification of an SPSD matrix, and therefore is itself SPSD. From (39a) again, note that we have

$$\begin{aligned} \begin{bmatrix} Q_k & M_k \\ M_k^T & R_k \end{bmatrix} &= \begin{bmatrix} Q & M \\ M^T & R \end{bmatrix} + \begin{bmatrix} D^T \\ -G^T \end{bmatrix} (\Sigma_k^D)^{-1} [D, -G] \\ &\quad + \begin{bmatrix} H^T [(\Sigma_k^H)^{-1} - (\Sigma_k^H Z_k \Sigma_k^H)^{-1}] H & 0 \\ 0 & 0 \end{bmatrix}, \end{aligned}$$

for all $k=1, 2, \dots, N-1$. Because of (44), we have that the left-hand side of this expression is a sum of SPSD terms, and therefore is itself SPSD. Finally, from (39a) note that each R_k , $k=0, 1, \dots, N-1$, is the sum of a

rows from (45) to obtain the following subsystem:

$$\begin{bmatrix} -I & Q_{k-1} & M_{k-1} & A^T \\ & M_{k-1}^T & R_{k-1} & B^T \\ & A & B & 0 \\ & & & -I & \Pi_k \end{bmatrix} \begin{bmatrix} \Delta \hat{p}_{k-2} \\ \Delta \hat{x}_{k-1} \\ \Delta \hat{u}_{k-1} \\ \Delta \hat{p}_{k-1} \\ \Delta \hat{x}_k \end{bmatrix} = \begin{bmatrix} \tilde{r}_{k-1}^x \\ \tilde{r}_{k-1}^u \\ \tilde{r}_{k-1}^p \\ \pi_k \end{bmatrix}. \quad (48)$$

Elimination of $\Delta \hat{p}_{k-1}$ and $\Delta \hat{x}_k$ yields

$$\begin{bmatrix} -I & Q_{k-1} + A^T \Pi_k A & A^T \Pi_k B + M_{k-1} \\ 0 & B^T \Pi_k A + M_{k-1}^T & R_{k-1} + B^T \Pi_k B \end{bmatrix} \begin{bmatrix} \Delta \hat{p}_{k-2} \\ \Delta \hat{x}_{k-1} \\ \Delta \hat{u}_{k-1} \end{bmatrix} = \begin{bmatrix} \tilde{r}_{k-1}^x + A^T \Pi_k \tilde{r}_{k-1}^p + A^T \pi_k \\ \tilde{r}_{k-1}^u + B^T \Pi_k \tilde{r}_{k-1}^p + B^T \pi_k \end{bmatrix}. \quad (49)$$

Finally, elimination of $\Delta \hat{u}_{k-1}$ yields the equation

$$-\Delta \hat{p}_{k-2} + \Pi_{k-1} \Delta \hat{x}_{k-1} = \pi_{k-1}, \quad (50)$$

where

$$\begin{aligned} \Pi_{k-1} &= Q_{k-1} + A^T \Pi_k A \\ &\quad - (A^T \Pi_k B + M_{k-1})(R_{k-1} + B^T \Pi_k B)^{-1} (B^T \Pi_k A + M_{k-1}^T), \end{aligned} \quad (51a)$$

$$\begin{aligned} \pi_{k-1} &= \tilde{r}_{k-1}^x + A^T \Pi_k \tilde{r}_{k-1}^p + A^T \pi_k \\ &\quad - (A^T \Pi_k B + M_{k-1})(R_{k-1} + B^T \Pi_k B)^{-1} (\tilde{r}_{k-1}^u + B^T \Pi_k \tilde{r}_{k-1}^p + B^T \pi_k). \end{aligned} \quad (51b)$$

Equation (51a) is the famous discrete-time Riccati equation for time-varying weighting matrices.

The solution of (45) can now be obtained as follows. We first set Π_N and π_n using (47), and then apply (51a) to obtain Π_k and π_k for $k = N-1, N-2, \dots, 1$. Next, we combine (46) for $k=1$ with the first two rows of (45); we obtain

$$\begin{bmatrix} R_0 & B^T \\ B & -I \\ & -I & \Pi_1 \end{bmatrix} \begin{bmatrix} \Delta \hat{u}_0 \\ \Delta \hat{p}_0 \\ \Delta \hat{x}_1 \end{bmatrix} = \begin{bmatrix} \tilde{r}_0^u \\ \tilde{r}_0^p \\ \pi_1 \end{bmatrix}, \quad (52)$$

and solve this system for $\Delta\hat{u}_0$, $\Delta\hat{x}_1$, $\Delta\hat{p}_0$. Next, we obtain from (49) and (48) that

$$\begin{aligned}\Delta\hat{u}_k &= (R_k + B^T \Pi_{k+1} B)^{-1} [\tilde{r}_k^u + B^T \Pi_{k+1} \tilde{r}_k^p + B^T \pi_{k+1} - (B^T \Pi_{k+1} A + M_k^T) \Delta\hat{x}_k], \\ \Delta\hat{x}_{k+1} &= A \Delta\hat{x}_k + B \Delta\hat{u}_k, \quad k = 1, 2, \dots, N-1.\end{aligned}$$

Finally, the steps $\Delta\hat{p}_k$ for $k = N-1, N-2, \dots, 1$ can be recovered from (46). The computational cost of the entire process is $O(N(m+n)^3)$.

The question of stability of this approach is an important one. Essentially the block elimination/Riccati scheme just described places restrictions on the pivot sequence, that is, the order in which the elements of the matrix in (45) are eliminated. Note that pivoting for numerical stability can occur internally, during the factorization of $(R_{k-1} + B^T \Pi_k B)$ in (51a) and (51b) for $k = N, N-1, \dots, 2$. In other circumstances, pivot restrictions are well known to lead to numerical instability, which manifests itself by blowup of the intermediate quantities that arise during the factorization, by which we mean that the intermediate quantities become much larger than the original data of the problem. However, in the present case, stability can be established by the simple argument of the next few paragraphs.

The coefficient matrix in (45) becomes increasingly ill-conditioned near the solution. This feature results from wide variation among the elements of the diagonal matrices $\Sigma_k^D, \Sigma_k^\epsilon, \Sigma_k^H$ defined by (36) which, as we see from (34), can vary between $\Theta(\mu)$ and $\Theta(\mu^{-1})$, where the duality measure μ approaches zero as the iterates approach the solution. It follows from (39a) that Q_k , $k = 1, 2, \dots, N$, has its eigenvalues in the range $[0, \Theta(\mu^{-1})]$, while positive definiteness of R ensures that the eigenvalues of R_k , $k = 0, 1, \dots, N-1$, lie in an interval $[\Theta(1), \Theta(\mu^{-1})]$. Since we showed earlier that the matrices

$$\begin{bmatrix} Q_k & M_k \\ M_k^T & R_k \end{bmatrix}, \quad k = 1, 2, \dots, N-1,$$

are SPSD, we deduce from the comments just made that their eigenvalues too must lie in the range $[0, \Theta(\mu^{-1})]$.

We now show that blowup does not occur during computation of the Riccati matrices Π_k and that, in all cases, their eigenvalues lie in the range $[0, \Theta(\mu^{-1})]$. This is certainly true of the starting matrix Π_N defined by (47). For the remaining matrices defined by (51a), we assume that our assertion is true for Π_k for some k , and prove that it continues to hold for Π_{k-1} . Note that the matrix

$$\begin{bmatrix} Q_{k-1} & M_{k-1} \\ M_{k-1}^T & R_{k-1} \end{bmatrix} + \begin{bmatrix} A^T \\ B^T \end{bmatrix} \Pi_k [A, B], \quad (53)$$

has both terms SPSD, with eigenvalues in the range $[0, \Theta(\mu^{-1})]$. Since Π_{k-1} is the Schur complement of $R_{k-1} + B^T \Pi_k B$ in the matrix (53), it must be positive semidefinite. Note that Π_{k-1} is well defined by the formula (51a), since $R_{k-1} + B^T \Pi_k B$ is an SPSD modification of the SPD matrix R , and so its inverse is well defined. Moreover, we can see from (51a) that Π_k is obtained by subtracting an SPSD matrix from the SPSD matrix $Q_{k-1} + A^T \Pi_k A$, and so its eigenvalues are bounded above by the eigenvalues of the latter matrix. By combining these observations, we conclude that the eigenvalues of Π_{k-1} lie in the range $[0, \Theta(\mu^{-1})]$, as claimed.

For the vectors π_k , $k = N, N-1, \dots, 1$, we have from the invertibility of $R_{k-1} + B^T \Pi_k B$ that they are well defined. Moreover, since the smallest eigenvalue of $R_{k-1} + B^T \Pi_k B$ has size $\Theta(1)$, we have from the formula (51b) and the estimate $\|\Pi_k\| = O(\mu^{-1})$ from the previous paragraph that $\|\pi_k\| = O(\mu^{-2})$, and so this vector does not blow up with k either. In fact, a more refined analysis can be used to deduce that $\|\pi_k\| = O(\mu^{-1})$, but we omit the details of this argument here.

We conclude that numerical instability is not a problem in applying the block elimination/Riccati scheme and that, in fact, we can expect this scheme to be as stable as any general scheme based on Gaussian elimination with pivoting.

It might be expected that the inherent ill conditioning of the system (45) may lead to an inaccurate computed solution, even when our numerical scheme is stable. It has long been observed by interior-point practitioners, however, that the computed steps are surprisingly effective steps for the algorithm, even on later iterations on which μ is tiny. This observation has recently found some theoretical support [see Wright (Refs. 23–24)], but the issues involved are beyond the scope of this paper.

3.4. Block Elimination: Constrained Endpoint. When endpoint constraints are present in the problem, they can be accounted for by adding extra recursions to the scheme of the previous section. We describe this approach below, but first mention an alternative way to handle the problem. The presence of endpoint constraints in the model is often symptomatic of the transition matrix A having eigenvalues outside the unit circle. In these circumstances, it is known that Riccati-based techniques can encounter stability difficulties. These difficulties are ameliorated by the technique of parameterizing the input as

$$u_k = Lx_k + r_k,$$

where L is a linear stabilizing feedback gain for (A, B) , as mentioned in Section 2.1. Alternatively, we can simply discard the Riccati strategy and instead apply a standard banded Gaussian-elimination scheme with partial

so that $y_2 = \Delta\beta$ can be found directly by substituting into (55a). We recover the remainder of the solution vector from (55b) by noting that

$$T_{11}^{-1} r_1 - T_{11}^{-1} T_{12} y_2 = \begin{bmatrix} \Delta\hat{u}_0 \\ \Delta\hat{p}_0 \\ \Delta\hat{x}_1 \\ \Delta\hat{u}_1 \\ \Delta\hat{p}_1 \\ \Delta\hat{x}_2 \\ \vdots \\ \Delta\hat{x}_N \end{bmatrix} - \begin{bmatrix} \Phi_0^u \\ \Phi_0^p \\ \Phi_1^x \\ \Phi_1^u \\ \Phi_1^p \\ \Phi_2^x \\ \vdots \\ \Phi_N^x \end{bmatrix} \Delta\beta.$$

In the implementation, the recurrences for computing Π_k , Ψ_k , π_k take place simultaneously, as do the recurrences needed for solving the systems (45) and (56). The additional cost associated with the n_f endpoint constraints is $O(N(m+n)^2 n_f)$. When $n_f < n$, which is a necessary condition for (38) to have a unique solution, the cost of solving the full system (38) is less than double the cost of solving the subsystem (45) alone by the method of the preceding section.

3.5. Hot Starting. Model predictive control solves a sequence of similar optimal control problems in succession. If the model is accurate and the disturbances are modest, the solution of one optimal control problem can be shifted one time step forward to yield a good approximation to the solution of the next problem in the sequence. Unfortunately, an approximate solution of this type is not a suitable starting guess for the interior-point method, since it usually lies at the boundary of the feasible region, whereas interior-point methods prefer starting point that strictly satisfy the inequalities in the constraint set. Starting points close to the so-called central path are more suitable. In the notation of Section 3.1, the characteristics of such points are that their pairwise products $\lambda_i t_i$ are similar in value for $i = 1, 2, \dots, m$ and that the ratio of the KKT violations in (23a), measured by $\mathcal{F}(z, \pi, \lambda, t)$, to the duality gap μ is not too large. We can attempt to find near-central points by bumping components of the shifted starting point off their bound. In the notation of Section 3.1, we turn the zero value of either t_i or λ_i into a small positive value. A second technique is to use a shifted version of one of the earlier interior-point iterates from the previous problem. Since the interior-point algorithm tends to follow the central path, and since the central path is sensitive to data perturbations only near the solution, this strategy generally produces an iterate that is close to the central path for the new optimal control subproblem.

In the presence of new disturbances, the previous solution has little relevance to the new optimal control problem. A starting point can be constructed from the unconstrained solution, or we can perform a cold start from a well-centered point, as is done to good effect in linear programming codes; see Wright (Ref. 21).

4. Computational Results

To gauge the effectiveness of the structured interior-point approach, we tested it against the standard quadratic programming approach, in which the states x_k are eliminated from the problem (19), (20) by using the model equation (20b). A reduced problem with unknowns u_k , $k=0, 1, \dots, N-1$, and ϵ_k , $k=1, 2, \dots, n$, is obtained. The reduction in dimension is accompanied by filling in the constraint matrices and the Hessian of the objective. The resulting problem is solved with the widely used code QPSOL (Ref. 22), which implements an active set method using dense linear algebra calculations.

We compared these two approaches on three common applications of the model predictive control methodology.

Example 4.1. Copolymerization Reactor. Congalidis et al. (Ref. 25) presented the following normalized model for the copolymerization of methyl methacrylate (MMA) and vinyl acetate (VA) in a continuous stirred tank reactor:

$$G(s) = \begin{bmatrix} \frac{0.34}{0.85s+1} & \frac{0.21}{0.42s+1} & \frac{0.50(0.50s+1)}{12s^2+0.4s+1} & 0 & \frac{6.46(0.9s+1)}{0.07s^2+0.3s+1} \\ -\frac{0.41}{2.41s+1} & \frac{0.66}{1.51s+1} & \frac{-0.3}{1.45s+1} & 0 & \frac{-3.72}{0.8s+1} \\ \frac{0.30}{2.54s+1} & \frac{0.49}{1.54s+1} & \frac{-0.71}{1.35s+1} & \frac{-0.20}{2.71s+1} & \frac{-4.71}{0.008s^2+0.41s+1} \\ 0 & 0 & 0 & 0 & \frac{1.02}{0.07s^2+0.31s+1} \end{bmatrix}$$

The normalized inputs into the system are the flows of monomer MMA (u_1), monomer VA (u_2), initiator (u_3), transfer agent (u_4), and temperature of the reactor jacket (u_5). The normalized outputs of the systems are the polymer production rate (y_1), mole fraction of MMA in the polymer (y_2), average molecular weight of the polymer (y_3), and reactor temperature (y_4). The model was realized in block observer canonical form (Ref. 27), where the dimension n of state after the realization is 18 and the number m of inputs is 5. The model was discretized with a sample period of 1.

The normalized inputs were constrained to be within 10% of their nominal operating steady-state values. The tuning parameters were chosen to be $Q = C^T C$ (where C is the measurement matrix obtained from the state space realization), while $M = 0$, $R = (0.1)I$, and the number of stages N is 100. Due to the very slow dynamics of the reactor, \bar{Q} was obtained from the solution of (8). The parameters z and Z are vacuous, since there are no soft constraints on the state. The controller was simulated with the following state disturbance:

$$[x_0]_j = 0.02 \sin j.$$

The interior-point method required 14 iterations to solve the optimization problem. Figure 1 shows the optimal control profile normalized with the upper bounds on the input constraints.

Example 4.2. Gage Control of a Polymer Film Process. We considered the gage (cross-directional control) of a 26-lane polymer film process with 26 actuators. We used the following model for our simulation:

$$A = 0.9I, \quad B = (I - A)K,$$

where the steady-state gain matrix K was extrapolated from data obtained from a 3M polymer film pilot plant. For this example, the dimension of the state n is 26 and the number m of inputs is 26. The state $[x]_j$ denotes the deviated film thickness in the j th lane and the input $[u]_j$ denotes the deviated position of the j th actuator.

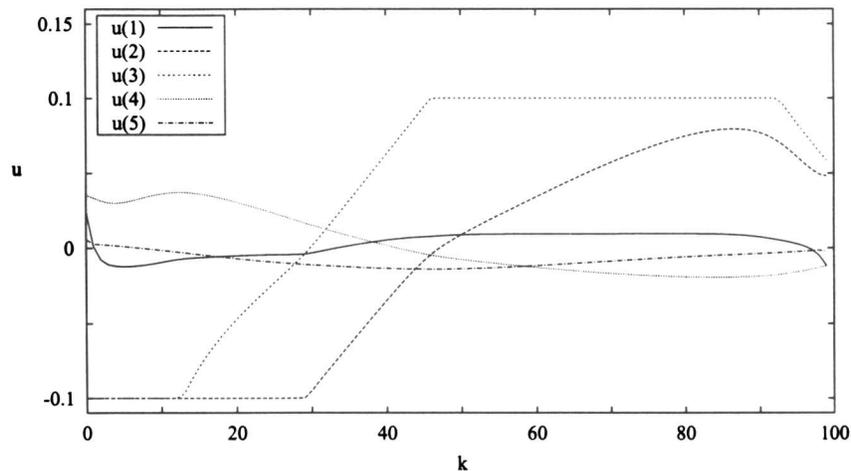


Fig. 1. Input profile for Example 4.1.

The actuators were constrained between the values of 0.1 and -0.1 , while the velocity of the actuators was constrained between the values of 0.025 and -0.025 . Since a large difference between actuator positions can create excess stress on the die, we imposed the following restriction on the change in input from stage to stage:

$$|[u]_j - [u]_{j-1}| < 0.05, \quad j = 2, 3, \dots, m.$$

We chose the tuning parameters to be

$$Q = I, \quad R = I, \quad S = I.$$

The matrix \bar{Q} was obtained from the solution of (10). The parameters z and Z are vacuous, since there are no soft constraints on the state. We chose a horizon of $N = 30$ to guarantee that the constraints were satisfied on the infinite horizon. The interior-point method required 11 iterations. Figure 2 shows the calculated optimal input profiles.

Example 4.3. Evaporator. Ricker et al. (Ref. 26) presented the following model for an evaporation process in a kraft pulp mill:

$$G(s) = \begin{bmatrix} \frac{1}{30s+1} & 0 \\ \frac{648s}{(30s+1)(20s+1)} & \frac{2.7(-6s+1)}{(20s+1)(5s+1)} \\ \frac{-90s}{(30s+1)(30s+1)} & \frac{-0.1375(-4s+1)}{(30s+1)(2.6s+1)} \end{bmatrix}.$$

The normalized outputs of the process are the feed level (y_1), product concentration (y_2), and product level (y_3). The normalized inputs for the process

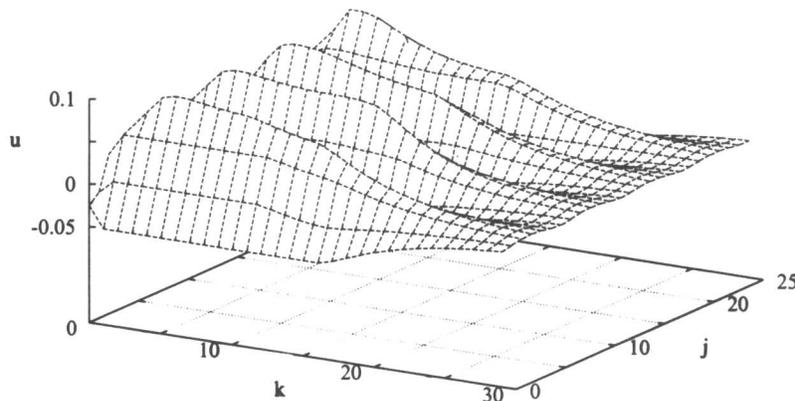


Fig. 2. Input profile for Example 4.2.

are the feed level setpoint (u_1) and steam flow (u_2). The process was realized in block observer canonical form (Ref. 27) and sampled every 0.5 minutes. The dimension n of the state after the realization is 9 and the number m of input is 3.

Both inputs were constrained to lie in the range $[-0.2, 0.2]$, while the three outputs were constrained to lie in $[-0.05, 0.05]$. A bound of 0.05 was also imposed on the input velocity. The controller was tuned with

$$Q = C^T C, \quad R = I, \quad Z = 0, \quad N = 60.$$

The matrix \bar{Q} was obtained from the solution of (10). A constant l_1 -penalty of 1000 was sufficient to force the soft constraints to hold when the solution is feasible. We simulated the controller with the following state disturbance:

$$[x_0]_j = \sin j + \cos j.$$

The interior-point method required 18 iterations to solve the optimization problem. Figure 3 shows the calculated optimal input profile, while Fig. 4 shows the predicted output profile. Note that the constraints for y_2 and y_3 are initially violated. The constraint for y_2 is feasible when $k \geq 8$ and the constraint for y_3 is feasible when $k \geq 34$. Increasing the l_1 -penalty did not change the resulting solution. Decreasing the l_1 -penalty leads to less aggressive control action, but the constraints are violated for a longer duration.

The computational times required by the structured interior-point approach and the naive quadratic programming approach are shown in

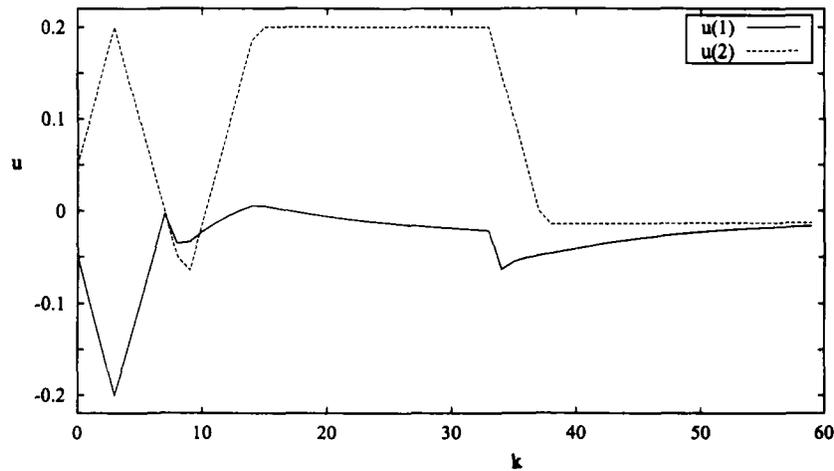


Fig. 3. Input profile at $t=0$ for Example 4.3.

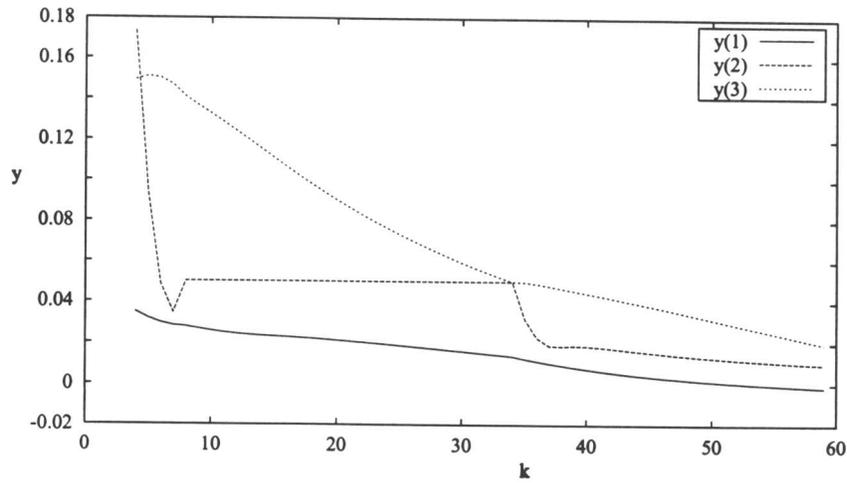


Fig. 4. Predicted output profile for Example 4.3.

Table 1. Our platform was a DEC Alphastation 250, and the times were obtained with the Unix *time* command. We used the value $\gamma=0.995$ in (29) as the proportion of maximum step to the boundary taken by our algorithm.

For the chosen (large) values of the horizon parameter N , the structured interior-point method easily outperforms the naive quadratic programming approach. For the latter approach, we do not include the time required to eliminate the states. These times were often quite significant, but they are calculated offline. For small values of the horizon parameter N , the naive quadratic programming approach outperforms the structured interior-point method, since the bandwidth is roughly the same relative order of magnitude as the dimensions of (38).

5. Concluding Remarks

We conclude with four brief comments on the structured interior-point method for MPC.

Table 1. Computational times (sec).

Example	Structured interior-point	Naive quadratic programming
4.1	3.80	23.78
4.2	20.33	276.91
4.3	2.01	25.32

(i) The structured method presented is also directly applicable to the dual problem of MPC, the constrained moving horizon estimation problem. In fact, the estimation problem will provide greater justification for structured approach because long horizons N arise frequently in this context. However, we did not investigate applying the structured optimization approach because the theory for linear constrained receding-horizon estimators is still in its infancy.

(ii) We can extend the structured method to nonlinear MPC by applying the approach of this paper to the linear-quadratic subproblems generated by sequential quadratic programming. Wright (Ref. 7), Arnold et al. (Ref. 6), and Steinbach (Ref. 9) all apply a similar technique to discrete-time optimal-control problems. While some theory for nonlinear MPC is available, the questions of robust implementation and suitable formulation of nonlinear MPC have not been resolved. See Mayne (Ref. 28) for a discussion of some of the issues.

(iii) Since the computational cost of the proposed algorithm is $O(N(m+n)^3)$, systems with large numbers of states and inputs can still present formidable computational challenges. Since large systems tend to be sparse (that is, A and B tend to be sparse, while Q and R tend to be nearly diagonal), we expect substantial increases in computational performance by exploiting the sparsity in (38) through the use of sparse matrix solvers. Since the sparsity tends to be structured in many applications, different strategies are preferable for different classes of processes. See, for example, the paper of Rao et al. (Ref. 29), who investigated strategies for further decomposing the problem structure in the gage control of sheet and film forming processes.

(iv) The last comment concerns time delays, which occur when more than one sampling period elapses before an input u_k affects the state of the system. In the simplest case, we can rewrite the state equation (5a) as

$$x_{k+1} = Ax_k + Bu_{k-d}, \quad (58)$$

for the case in which the delay is d sampling periods. The natural infinite-horizon LQR objective function for this case is

$$\Phi = (1/2) \sum_{k=0}^{\infty} (x_k^T Q x_k + u_k^T R u_k + 2x_{k+d}^T M u_k), \quad (59)$$

where the cross-penalty term relates u_{k-d} and x_k . Since the first $d+1$ state vectors x_0, x_1, \dots, x_d are independent of the inputs, the decision variables in the optimization problem are x_{d+1}, x_{d+2}, \dots and u_0, u_1, \dots . By defining

$$\tilde{x}_k = x_{k+d}, \quad k=0, 1, 2, \dots,$$

and removing constant terms from (59), the objective function and state equation become

$$\Phi' = (1/2) \sum_{k=0}^{\infty} (\tilde{x}_k^T Q \tilde{x}_k + u_k^T R u_k + 2\tilde{x}_k^T M u_k), \tag{60}$$

$$\tilde{x}_0 = x_d, \quad \tilde{x}_{k+1} = A \tilde{x}_k + B u_k, \quad k=0, 1, 2, \dots \tag{61}$$

These formulas have the same form as (4)–(5).

If no additional constraints of the form (5b) are present, a Riccati equation may be used to solve (60)–(61) directly, as in Section 2.2. If state constraints of the form $Hx_k \leq h$ or jump constraints of the form $G\Delta u_k \leq g$ are present [as in (1)], we can still apply constraint softening (Section 2.3) and use the approaches described in Section 2.2. To obtain finite-horizon versions of (60)–(61), the techniques of Section 3.1 can then be used to solve the problem efficiently.

Difficulties may arise when multiple time delays are present, since these may reduce the locality of the relationships between the decision variables and lead to significant broadening of the bandwidths of the matrices in (35) and (38). A process in which two time delays are present, with d_1 and d_2 sampling intervals, can be described by a state equation of the following form:

$$\begin{bmatrix} x_{k+1}^1 \\ x_{k+1}^2 \end{bmatrix} = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} x_k^1 \\ x_k^2 \end{bmatrix} + \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix} \begin{bmatrix} u_{k-d_1}^1 \\ u_{k-d_2}^2 \end{bmatrix}.$$

A problem with this dynamics can be solved by augmenting the state vector x_k with the input variables $u_{k-d_1}, u_{k-d_1-1}, \dots, u_{k-d_2+1}$ (assuming that $d_2 > d_1$) and applying the technique for a single time delay outlined above. Alternatively, the KKT conditions for the original formulation can be used directly as the basis of an interior-point method. The linear system to be solved at each interior-point iteration will contain not only diagonal blocks of the form in (35), but also a number of blocks at some distance from the diagonal. Some rearrangement to reduce the overall bandwidth may be possible, but expansion of the bandwidth by an amount proportional to $(d_2 - d_1)m$ is inevitable.

Of course, we can also revert to the original approach of eliminating the states x_0, x_1, \dots from the problem to obtain a problem in which the inputs u_0, u_1, \dots alone are decision variables. The cost of this approach, too, is higher than in the no-delay case, because the horizon length N usually must be increased to incorporate the effects of the delayed dynamics. One could postulate that certain processes would be effectively handled by the standard approach while others would be effectively handled by the structured approach. Perhaps, the only solution is to exercise engineering

judgment to decompose the full control problem into smaller problems without large delays and treat the neglected delays connecting the decomposed systems as disturbances. This issue remains unresolved and is a topic of current research.

References

1. CHMIELEWSKI, D., and MANOUSIOUTHAKIS, V., *On Constrained Infinite-Time Linear-Quadratic Optimal Control*, System and Control Letters, Vol. 29, pp. 121–129, 1996.
2. SCOKAERT, P. O., and RAWLINGS, J. B., *Constrained Linear-Quadratic Regulation*, IEEE Transactions on Automatic Control, Vol. 43, pp. 1163–1169, 1998.
3. SZNAIER, M., and DAMBORG, M. J., *Suboptimal Control of Linear Systems with State and Control Inequality Constraints*, Proceedings of the 26th Conference on Decision and Control, pp. 761–762, 1987.
4. BERTSEKAS, D. P., *Dynamic Programming*, Prentice-Hall, Englewood Cliffs, New Jersey, 1987.
5. GLAD, T., and JOHNSON, H., *A Method for State and Control Constrained Linear-Quadratic Control Problems*, Proceedings of the 9th IFAC World Congress, Budapest, Hungary, pp. 1583–1587, 1984.
6. ARNOLD, E., TATJEWSKI, P., and WOLOCZOWICZ, P., *Two Methods for Large-Scale Nonlinear Optimization and Their Comparison on a Case Study of Hydro-power Optimization*, Journal of Optimization Theory and Applications, Vol. 81, pp. 221–248, 1994.
7. WRIGHT, S. J., *Interior-Point Methods for Optimal Control of Discrete-Time Systems*, Journal of Optimization Theory and Applications, Vol. 77, pp. 161–187, 1993.
8. WRIGHT, S. J., *Applying New Optimization Algorithms to Model Predictive Control*, Chemical Process Control-V, AIChE Symposium Series, Vol. 93, pp. 147–155, 1997.
9. STEINBACH, M. C., *A Structured Interior-Point SQP Method for Nonlinear Optimal Control Problems*, Computational Optimal Control, Edited by R. Burlirsch and D. Kraft, Birkhäuser, Basel, Switzerland, pp. 213–222, 1994.
10. LIM, A., MOORE, J., and FAYBUSOVICH, L., *Linearly Constrained LQ and LQG Optimal Control*, Proceedings of the 13th IFAC World Congress, San Francisco, California, 1996.
11. MEHROTRA, S., *On the Implementation of a Primal-Dual Interior-Point Method*, SIAM Journal on Optimization, Vol. 2, pp. 575–601, 1992.
12. MUSKE, K. R., and RAWLINGS, J. B., *Model Predictive Control with Linear Models*, AIChE Journal, Vol. 39, pp. 262–287, 1993.
13. SCOKAERT, P. O., and RAWLINGS, J. B., *Infinite-Horizon Linear-Quadratic Control with Constraints*, Proceedings of the 13th IFAC World Congress, San Francisco, California, pp. 109–113, 1996.

14. KEERTHI, S. S., *Optimal Feedback Control of Discrete-Time Systems with State-Control Constraints and General Cost Functions*, PhD Thesis, University of Michigan, 1986.
15. ROSSITER, J. A., RICE, M. J., and KOUVARITAKIS, B., *A Robust Stable State-Space Approach to Stable Predictive Control Strategies*, Proceedings of the American Control Conference, Albuquerque, New Mexico, pp. 1640-1641, 1997.
16. RAWLINGS, J. B., and MUSKE, K. R., *Stability of Constrained Receding-Horizon Control*, IEEE Transactions on Automatic Control, Vol. 38, pp. 1512-1516, 1993.
17. MEADOWS, E. S., MUSKE, K. R., and RAWLINGS, J. B., *Implementable Model Predictive Control in the State Space*, Proceedings of the 1995 American Control Conference, pp. 3699-3703, 1995.
18. GILBERT, E. G., and TAN, K. T., *Linear Systems with State and Control Constraints: The Theory and Application of Maximal Output Admissible Sets*, IEEE Transactions on Automatic Control, Vol. 36, pp. 1008-1020, 1991.
19. FLETCHER, R., *Practical Methods of Optimization*, John Wiley and Sons, New York, New York, 1987.
20. HAGER, W. W., *Lipschitz Continuity for Constrained Processes*, SIAM Journal on Control and Optimization, Vol. 17, pp. 321-338, 1979.
21. WRIGHT, S. J., *Primal-Dual Interior-Point Methods*, SIAM Publications, Philadelphia, Pennsylvania, 1997.
22. GILL, P. E., MURRAY, W., SAUNDERS, M. A., and WRIGHT, M. H., *User's Guide for SOL/QPSOL: A Fortran Package for Quadratic Programming*, Technical Report SOL 83-12, Systems Optimization Laboratory, Department of Operations Research, Stanford University, 1983.
23. WRIGHT, S. J., *Modified Cholesky Factorizations in Interior-Point Algorithms for Linear Programming*, Preprint ANL/MCS-P600-0596, Mathematics and Computer Science Division, Argonne National Laboratory, 1996.
24. WRIGHT, S. J., *Stability of Augmented System Factorizations in Interior-Point Methods*, SIAM Journal on Matrix Analysis and Its Applications, Vol. 18, pp. 191-222, 1997.
25. CONGALIDIS, J. B., RICHARDS, J. R., and RAY, W. H., *Modeling and Control of a Copolymerization Reactor*, Proceedings of the American Control Conference, Seattle, Washington, pp. 1779-1793, 1986.
26. RICKER, N. L., SUBRAMANIAN, T., and SIM, T., *Case Studies of Model-Predictive Control in Pulp and Paper Production*, Proceedings of the 1988 IFAC Workshop on Model-Based Process Control, Edited by T. J. McAvoy, Y. Arkun, and E. Zafiriou, Pergamon Press, Oxford, England, pp. 13-22, 1988.
27. CHEN, C. T., *Linear System Theory and Design*, Holt, Rhinehart, and Winston, New York, New York, 1984.
28. MAYNE, D. Q., *Nonlinear Model Predictive Control: An Assessment*, Chemical Process Control-V, AIChE Symposium Series, Vol. 93, pp. 217-231, 1997.
29. RAO, C. V., CAMPBELL, J. C., RAWLINGS, J. B., and WRIGHT, S. J., *Efficient Implementation of Model Predictive Control for Sheet and Film Forming Processes*, Proceedings of American Control Conference, Albuquerque, New Mexico, pp. 2940-2944, 1997.

Reproduced with permission of the copyright owner. Further reproduction prohibited without permission.