

for $\omega \approx 3.4868$. We conclude that the minimal volume is $v_{\min} \approx 8.45$ and the relative volume is

$$\frac{v_{\min}}{v(r^-)} \approx 11.2667.$$

VI. CONCLUSION

As an alternative to the classical robustness margin, in this paper we introduced the notion of minimum destabilizing volume. We have shown that this volume can be easily computed if the control system is affected by interval and affine parametric uncertainties. A problem of subsequent interest is the extension of these results to more general uncertainty structures.

REFERENCES

- [1] J. Ackermann, A. Bartlett, D. Kaesbauer, W. Sienel, and R. Steinhauser, *Robust Control, Systems with Uncertain Physical Parameters*. London, U.K.: Springer-Verlag, 1993.
- [2] B. R. Barmish, *New Tools for Robustness of Linear Systems*. New York: Macmillan, 1994.
- [3] B. R. Barmish, P. P. Khargonekar, Z. C. Shi, and R. Tempo, "Robustness margin need not be a continuous function of the problem data," *Syst. Contr. Lett.*, vol. 15, pp. 91–98, 1990.
- [4] B. R. Barmish and B. Polyak, "The volumetric singular value and robustness of feedback control systems," in *Proc. 32nd Conf. Decision and Control*, San Antonio, TX, 1993, pp. 521–522.
- [5] A. C. Bartlett, C. V. Hollot, and L. Huang, "Root locations of an entire polytope of polynomials: It suffices to check the edges," *Math. Contr., Signals, and Syst.*, vol. 1, pp. 61–71, 1988.
- [6] S. P. Bhattacharyya, H. Chapellat, and L. H. Keel, *Robust Control: The Parametric Approach*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [7] T. E. Djaferis, *Robust Control Design: A Polynomial Approach*. Boston, MA: Kluwer, 1995.
- [8] T. E. Djaferis and C. V. Hollot, "The stability of a family of polynomials can be deduced from a finite number $O(k^3)$ of frequency checks," *IEEE Trans. Automat. Contr.*, vol. 34, pp. 982–986, 1989.
- [9] R. Horst and H. Tui, "Global optimization: A deterministic approach," in *Control and Information Sciences*. Berlin, Germany: Springer-Verlag, 1996.
- [10] V. L. Kharitonov, "Asymptotic stability of an equilibrium position of a family of systems of linear differential equations," *Differentsial'nye Uravneniya*, vol. 14, pp. 1483–1485, 1978.
- [11] J. Kogan, "Robust stability and convexity," in *Control and Information Sciences*. Berlin, Germany: Springer-Verlag, 1995.
- [12] M. Marden, *Geometry of Polynomials, Mathematical Surveys*, Providence, RI: Amer. Math. Soc., no. 3.

Constrained Linear Quadratic Regulation

Pierre O. M. Sokaert and James B. Rawlings

Abstract—This paper is a contribution to the theory of the infinite-horizon linear quadratic regulator (LQR) problem subject to inequality constraints on the inputs and states, extending an approach first proposed by Szaier and Damborg [16]. A solution algorithm is presented, which requires solving a finite number of finite-dimensional positive definite quadratic programs. The constrained LQR outlined does not feature the undesirable mismatch between open-loop and closed-loop nominal system trajectories, which is present in the other popular forms of model predictive control (MPC) that can be implemented with a finite quadratic programming algorithm. The constrained LQR is shown to be both optimal and stabilizing. The solution algorithm is guaranteed to terminate in finite time with a computational cost that has a reasonable upper bound compared to the minimal cost for computing the optimal solution. Inherent to the approach is the removal of a tuning parameter, the control horizon, which is present in other MPC approaches and for which no reliable tuning guidelines are available. Two examples are presented that compare constrained LQR and two other popular forms of MPC. The examples demonstrate that constrained LQR achieves significantly better performance than the other forms of MPC on some plants, and the computational cost is not prohibitive for online implementation.

Index Terms—Constraints, infinite horizon, linear quadratic regulation, model predictive control.

I. INTRODUCTION

In 1960, Kalman [4] showed that the Riccati equation associated with the finite-horizon linear quadratic regulator (LQR) has a well-defined limit and used that result to solve the infinite-horizon LQR problem. To date, this remains one of the most influential discoveries of the modern control era. In the late 1970's, Richalet *et al.* [12] and Cutler and Ramaker [3] emulated the finite-horizon LQR for constrained processes, marking the beginning of the industrial implementation of what comes to be known as model predictive control (MPC). In the 1980's, the theoretical development of MPC with constraints ran into serious difficulties, and it became increasingly apparent that a return to an infinite-horizon formulation is required to produce stabilizing control laws [1]. Various contributions are made in the early 1990's, introducing infinite horizons into the MPC framework for constrained linear processes [11]. However, rather than address the full infinite-horizon constrained LQR problem, all emerging MPC variants, except the one by Szaier and Damborg [16] discussed below, rely on a finite and suboptimal parameterization of the postulated control sequence. In this paper, the use of a finite-input parameterization is relaxed, leading to the formulation of a control scheme we call constrained LQR.

In their remarkably succinct paper, Szaier and Damborg [16] present the basics of the approach. They treat a more restricted class of problems than the one considered here. In their problem statement, the state and input are constrained to lie in bounded convex polyhedrons. We make no restriction on boundedness of the constraint region. A major concern in their work is the real-time implementation limits on calculation time, which can cause early termination of their algorithm prior to optimal solution. In this work we have found that

Manuscript received December 13, 1995. This work was supported in part by the National Science Foundation under Grant CTS-9311420 and by the Texas Advanced Technology Program under Grant 003658-078.

The authors are with the Department of Chemical Engineering, University of Wisconsin, Madison, Madison, WI 53706 USA.

Publisher Item Identifier S 0018-9286(98)04634-0.

for many examples, the cost of our algorithm is not much greater and is sometimes less than the cost of the standard suboptimal MPC algorithms already in use in industrial process control applications. Therefore, many industrial applications exist for which this approach can be implemented as an improvement over current practice. The constrained LQR approach also removes what we consider in current MPC approaches to be a nuisance tuning parameter, the control horizon N , i.e., the number of future control moves considered in the optimization.

Sznaier and Damborg do not provide a full discussion of the conditions under which their algorithm is stabilizing nor prove that it terminates in finite time. The stability of their controller is not guaranteed if their algorithm terminates before the optimal solution is calculated. In Sznaier's thesis [15], a contraction constraint is introduced, which, if feasible, does provide asymptotic stability; however, this added constraint leads to suboptimal control, even when ample computational time is available. We provide more complete results on termination and stability and also highlight the benefits of constrained LQR compared to other MPC variants in current use.

The paper is organized as follows. The constrained LQR law is defined in Section II, and its stabilizing properties are established. Section III outlines the details of a practical implementation of constrained LQR, which is shown to require the solution of a finite number of finite-dimensional positive definite quadratic programs. In Section IV, we discuss the computational aspects of the constrained LQR algorithm and show that the computational cost has a reasonable upper bound, compared to the minimal cost for computing the optimal solution. Finally, examples are presented in Section V, where it is demonstrated that constrained LQR achieves significantly better performance than other forms of MPC on some plants, and the computational cost is not prohibitive for online implementation. Concluding remarks are made in Section VI.

II. OVERVIEW OF THE PROBLEM STATEMENT

We consider time-invariant, linear, and discrete-time systems described by the state-space equation

$$x_{t+1} = Ax_t + Bu_t \quad (1)$$

in which $x_t \in \mathbb{R}^n$ and $u_t \in \mathbb{R}^m$ are the state and input vectors at discrete time t and A and B are the state transition and input distribution matrices, respectively. It is assumed throughout the paper that (A, B) is stabilizable.

The control objective is to regulate the state of the system optimally to the origin. Optimality is defined in terms of a quadratic objective and a set of inequality constraints.

The objective is defined over an infinite horizon and is given by

$$\phi(x_t, \pi) = \sum_{j=t}^{\infty} x'_{j|t} Q x_{j|t} + u'_{j|t} R u_{j|t} \quad (2)$$

in which $Q \geq 0$ and $R > 0$ are symmetric weighting matrices, such that $(Q^{1/2}, A)$ is detectable, and

$$\pi = \{u_{t|t}, u_{t+1|t}, \dots\} \quad (3)$$

$$x_{j+1|t} = Ax_{j|t} + Bu_{j|t}, \quad t \leq j \quad (4)$$

with $x_{t|t} = x_t$. The constraints are also defined on an infinite horizon and take the form

$$\begin{aligned} Hx_{j+1|t} &\leq h, & t \leq j \\ Du_{j|t} &\leq d \end{aligned} \quad (5)$$

where $h \in \mathbb{R}_+^{n_h}$ and $d \in \mathbb{R}_+^{n_d}$ define the constraint levels, with n_h and n_d denoting the number of state and input constraints respectively, and H and D are the state and input constraint distribution matrices.

We summarize the three relevant control problems of interest in this paper.

Problem 0—LQR:

$$\begin{aligned} &\min_{\pi} \phi(x_t, \pi) \\ &\text{subject to:} \\ &x_{j+1|t} = Ax_{j|t} + Bu_{j|t}, \quad t \leq j. \end{aligned} \quad (6)$$

Problem 0 was formulated and solved by Kalman, and the solution is the well-known linear feedback control law

$$u_t = -Kx_t \quad (7)$$

in which the controller gain K can be calculated from the solution of the discrete algebraic Riccati equation. The linear feedback law of (7) is stabilizing under the above assumptions of stabilizability and detectability.

Problem 1—Constrained LQR:

$$\begin{aligned} &\min_{\pi} \phi(x_t, \pi) \\ &\text{subject to:} \\ &x_{j+1|t} = Ax_{j|t} + Bu_{j|t} \\ &Hx_{j+1|t} \leq h, \quad t \leq j \\ &Du_{j|t} \leq d. \end{aligned} \quad (8)$$

Problem 1 is a natural extension of the infinite-horizon LQR (Problem 0) that includes constraints. The only difficulty is the infinite number of decision variables in the optimization and the infinite number of constraints.

Problem 2—An MPC Problem:

$$\begin{aligned} &\min_{\pi} \phi(x_t, \pi) \\ &\text{subject to:} \\ &x_{j+1|t} = Ax_{j|t} + Bu_{j|t}, \quad t \leq j \\ &Hx_{j+1|t} \leq h, \\ &t \leq j \leq t + N - 1 \\ &Du_{j|t} \leq d, \\ &u_{j|t} = -Kx_{j|t}, \quad t + N \leq j \end{aligned} \quad (9)$$

This form of MPC has a finite number of decision variables, N , and a finite number of constraints, $N(n_d + n_h)$; it can therefore be solved with standard quadratic programming methods. The only interest in this problem here is as an aid in solving Problem 1. The idea of appending the $u = -Kx$ unconstrained linear law to the finite set of N decision variables is used in Michalska and Mayne's dual-mode controller for nonlinear systems [8]; see also [10] for a brief review.

Concentrating on Problem 1, the open-loop optimal control is obtained by minimization of the objective over all control profiles that satisfy the inequality constraints. The feedback law is then defined by receding horizon implementation of the optimal open-loop control. Given the open-loop optimal strategy $\pi^*(x_t) = \{u_{t|t}^*(x_t), u_{t+1|t}^*(x_t), \dots\}$, we therefore have the control law

$$u_t = g(x_t) \equiv u_{t|t}^*(x_t) \quad (10)$$

in which $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ denotes the nonlinear map between state and control.

Remark 1: For notational simplicity, we drop the x_t argument in $u_{j|t}^*(x_t)$ in the sequel, i.e., $u_{j|t}^* \equiv u_{j|t}^*(x_t)$. \square

A. Stabilizing Properties

Not surprisingly, the constrained LQR law (10) benefits from similar stabilizing properties as recent formulations of MPC.

Theorem 1 (Nominal Stability): The origin is an exponentially stable fixed point of the closed-loop system $x_{t+1} = Ax_t + Bg(x_t)$, with a region of attraction equal to the domain of g .

Proof: The proof proceeds by showing that, under the assumption that x_t belongs to the domain of g [i.e., there exists a control profile π that satisfies (4) and (5)], the optimal value of the objective, $\phi^*(x)$, is finite. This implies that $Q^{1/2}x_{j|t}, u_{j|t} \rightarrow 0$, as $j \rightarrow \infty$. Because nominal closed-loop performance is identical to the open-loop predictions, it follows that $Q^{1/2}x_t, u_t \rightarrow 0$, as $t \rightarrow \infty$; as $(Q^{1/2}, A)$ is detectable, this implies that $x_t \rightarrow 0$. Note that the assumption of stabilizability ensures that the domain of g includes at least an open neighborhood of the origin. ■

Moreover, in view of [13], we also have a perturbed stability guarantee.

Theorem 2 (Perturbed Stability): The origin is a locally asymptotically stable fixed point of the perturbed closed-loop system $x_{t+1} = Ax_t + Bg(x_t) + p_t$, if $\lim_{t \rightarrow \infty} p_t = 0$.

Proof: Under the assumption of feasibility, Muske [9] shows that $g(x)$ is Lipschitz continuous. As $g(x)$ is a nominally exponentially stabilizing control law, the perturbed stability result of the theorem follows directly from [13]. ■

Remark 2: The result of Theorem 2 is local; this means that p_t must be small enough, for all t , such that the trajectory of the perturbed system remains in the domain of g , i.e., the perturbation must not cause infeasibility. □

Remark 3: Note also that, if the perturbation p_t converges to zero exponentially, the result of Theorem 2 can be strengthened to exponential stability. □

This result is significant because it leads directly to a stability proof for the cascade of a stable state observer and the constrained LQR law.

III. IMPLEMENTATION

A. Notation

As in [16], we make the following definition.

Definition 1: $\mathbb{X}_K \subseteq \mathbb{R}^n$ denotes the set of states x_t for which the unconstrained LQR law, $u_{j|t} = -Kx_{j|t}$ ($t \leq j$), satisfies (4) and (5). □

It is immediately apparent that $\pi^*(x_t) = \{-Kx_{t|t}, -Kx_{t+1|t}, \dots\}$, for all $x_t \in \mathbb{X}_K$. Under the assumption that all elements of h and d are strictly positive, \mathbb{X}_K contains an open neighborhood of the origin [16].

Definition 2: B_r denotes a ball of radius $r > 0$, centered on the origin, such that $B_r \subseteq \mathbb{X}_K$. □

Definition 3: Given $N \in \mathbb{N}$ and the current value of the state x_t , $\mathbb{P}_N(x_t)$ denotes the set of control profiles π such that

$$\begin{aligned} Hx_{j+1|t} &\leq h, & t \leq j \leq t + N - 1 \\ Du_{j|t} &\leq d \\ u_{j|t} &= -Kx_{j|t}, & t + N \leq j. \end{aligned} \quad (11)$$

Further, we define $\mathbb{IP}(x_t)$ to be the set of control profiles that satisfy (4) and (5). $\mathbb{IP}(x_t)$ thus may be thought of as the limit of $\mathbb{P}_N(x_t)$ as $N \rightarrow \infty$. □

Definition 4: $\mathbb{W}_N \subseteq \mathbb{R}^n$ denotes the set of states x for which $\mathbb{P}_N(x) \neq \emptyset$. Similarly, $\mathbb{W} \subseteq \mathbb{R}^n$ is the set of states for which $\mathbb{P}(x) \neq \emptyset$, i.e., \mathbb{W} is the domain of g . □

B. Related Problem Statements

Using the notation defined above, the optimal open-loop constrained control problems described in Section II may be stated concisely as follows.

Problem 1—Constrained LQR: Given $x \in \mathbb{W}$, find a control profile $\pi^*(x)$ for which

$$\phi^*(x) = \min_{\pi \in \mathbb{P}} \phi(x, \pi) \quad (12)$$

is attained. □

Remark 4: Existence of $\pi^*(x)$ for all $x \in \mathbb{W}$ is easy to establish and follows from the results in the sequel. Uniqueness follows from the uniqueness of the minimum of a positive definite quadratic function on a convex set. □

A contribution of this paper is to show that under the usual assumption that all elements of h and d are strictly positive (which implies that no constraint passes through the origin and $r > 0$ —see Definition 2), the solution to Problem 1 can be calculated in finite computational time. We therefore show that suboptimality need not be a property of implementable constrained predictive control schemes.

We start by considering the related MPC problem in which the following finite input parameterization is employed:

$$u_{j|t} = -Kx_{j|t}, \quad j \geq t + N \quad (13)$$

where $N \in \mathbb{N}$ is the finite control horizon.

In this setting, the open-loop optimal control is obtained by solution of the familiar MPC optimization.

Problem 2—An MPC Problem: Given $N \in \mathbb{N}$, $x \in \mathbb{W}_N$, find the control profile $\pi_N(x) = \{u_{t|t}^N, u_{t+1|t}^N, \dots, -Kx_{t+N|t}^N, -Kx_{t+N+1|t}^N, \dots\}$ for which

$$\phi_N(x) = \min_{\pi \in \mathbb{P}_N} \phi(x, \pi) \quad (14)$$

is attained. □

This constrained minimization is an Nm -dimensional positive definite quadratic program with $N(n_c + n_d)$ constraints. For any finite horizon N , efficient solutions for Problem 2 are therefore easily formulated.

Remark 5: Existence of $\pi_N(x)$ follows, by definition, for all $x \in \mathbb{W}_N$. Uniqueness results from positive definiteness of the Hessian, which is guaranteed because R is positive definite. Note that in the context of Problem 2, we have

$$\phi(x, \pi) = \sum_{j=t}^{N-1} (x'_{j|t} Q x_{j|t} + u'_{j|t} R u_{j|t}) + x'_{t+N|t} \tilde{Q} x_{t+N|t} \quad (15)$$

where \tilde{Q} is the solution of the matrix Lyapunov equation

$$\tilde{Q} = Q + K' R K + (A - BK)' \tilde{Q} (A - BK). \quad (16)$$

C. Properties of ϕ^* , ϕ_N , π^* , and π_N

The interesting relationship between $\pi^*(x)$ and $\pi_N(x)$ is the motivation for our interest in Problem 2. We have the following results.

Definition 5: We denote by $x_{j|t}^*$ the state predictions that correspond to the optimal open-loop control profile $\pi^*(x_t)$ and by $x_{j|t}^N$ those that correspond to $\pi_N(x_t)$. □

Remark 6: Note that $x_{j|t}^*$ and $x_{j|t}^N$ are functions of x_t . That dependence is, however, left implicit for notational simplicity. □

Lemma 1: $x_{j|t}^* \in \mathbb{X}_K \iff u_{k|t}^* = -Kx_{k|t}^*, \forall k \geq j$.

Proof: The control profile that is optimal with respect to the constrained LQR objective is $u_{j|t} = -Kx_{j|t}$, for all $x_{j|t} \in \mathbb{X}_K$ —see Definition 1. In view of Bellman's principle of optimality, the sequence of controls $\{u_{t|t}^*, u_{t+1|t}^*, \dots\}$ is optimal over all time intervals $[k, \infty)$, $k \geq t$. The forward implication of the lemma therefore follows from optimality.

By definition, \mathbb{X}_K is the set of states outside which $u = -Kx$ violates the constraints of (4) and (5). As the optimal control profile $\{u_{j|t}^*, u_{j+1|t}^*, \dots\}$ satisfies these constraints, it follows that $u_{j|t}^* = -Kx_{j|t}^*$ only if $x_{j|t}^* \in \mathbb{X}_K$, which establishes the reverse implication of the lemma. ■

Theorem 3: For every $x \in \mathbb{W}$, there exists a finite $N_\infty(x) \in \mathbb{I}$, such that $\phi^*(x) = \phi_N(x)$, $\pi^*(x) = \pi_N(x)$, for all $N \geq N_\infty(x)$.

Proof: As \mathbb{X}_K contains an open neighborhood of the origin and $\pi^*(x)$ drives the state predictions $x_{j|t}^*$ to the origin, there exists a finite integer $N_\infty(x)$ such that $x_{t+N_\infty(x)|t}^* \in \mathbb{X}_K$. In view of Lemma 1, we therefore have for all $N \geq N_\infty(x)$, $\pi^*(x) \in \mathbb{I}P_N$, i.e., $\pi^*(x) \in \mathbb{I}P_N \cap \mathbb{I}P$. As $\pi_N(x)$ minimizes $\phi(x, \pi)$ in $\mathbb{I}P_N$, it follows that $\phi^*(x) = \phi_N(x)$, and consequently $\pi^*(x) = \pi_N(x)$ for all $N \geq N_\infty(x)$. ■

This result is key to our approach to solving Problem 1, as explained below.

D. The Control Algorithm

In view of the discussion in Section III-C, an algorithm that is guaranteed to identify the optimal control profile $\pi^*(x)$ in finite computational time is the following.

Algorithm 1—Constrained LQR:

Step 0) Choose a finite horizon N_0 , set $N = N_0$.

Step 1) Solve Problem 2.

Step 2) If $x_{t+N|t}^* \in \mathbb{X}_K$, go to Step 4).

Step 3) Increase N , go to Step 1).

Step 4) Terminate: $\pi^*(x) = \pi_N(x)$. □

Sznaier and [16] Damborg discuss a similar algorithm, although they do not show that termination occurs in a finite number of iterations. Finiteness of $N_\infty(x)$ in Theorem 3, however, ensures that termination occurs after a finite number of cycles, regardless of the choice of initial horizon in Step 0) and of the heuristic used to increase it in Step 3).

IV. COMPUTATIONAL ISSUES

The spirit of optimal control, which lies at the root of MPC and LQR, leaves no question that Problem 1 is a more natural formulation of the constrained predictive control law than other MPC formulations in current use. The solution of Problem 1 can be computationally taxing, however. It is therefore necessary to assess the computational demands of the method and it is desirable to minimize them. In this section, we discuss the computational aspects of Algorithm 1. We also propose a modification of Algorithm 1 that gives guaranteed stability, even when there is not sufficient time for the algorithm to terminate. Finally, we discuss upper bounds on $N_\infty(x)$ and how they may be used to obtain one-shot solutions to Problem 1.

A. Computational Aspects of Algorithm 1

The choice of the initial horizon, N_0 , in Step 0), has an effect on the efficiency of the algorithm, since an initial horizon that is close to or greater than $N_\infty(x)$ leads to early termination. However, $N_\infty(x)$ varies with x , and the effect of the initial horizon is generally not critical. An initial choice $N_0 = 0$ can therefore be recommended as the default, with the desirable side effect that no optimization needs to be performed close to steady state when the state belongs to \mathbb{X}_K .

The solution of Problem 2, in Step 1), is the computationally expensive part of Algorithm 1. Efficient quadratic program solutions should of course be implemented. Infeasible interior point methods appear to be well suited. On one hand, the computational price of that solution is $O(N)$ [18], c.f. active set methods for which it is $O(N^2)$. On the other hand, the starting point for the optimization can be infeasible and the quadratic program on one cycle of Algorithm 1

can therefore be warm-started from the solution at the previous cycle. As the solution of Problem 2 approaches that of Problem 1 when N increases, the computational burden of the quadratic program with warm-start should not rise dramatically as Algorithm 1 is cycled and N is increased.

Although the set \mathbb{X}_K may be expensive to calculate, the check that $x_{t+N|t}^* \in \mathbb{X}_K$, in Step 2), is inexpensive. In fact, the check is performed with no need for detailed information about \mathbb{X}_K . Predictions of the state and input are simply propagated until either a constraint violation is detected or the predicted state enters B_r . Note that this only requires prediction on a horizon that is guaranteed to be finite. The conclusion that $x_{t+N|t}^* \in \mathbb{X}_K$ is made if and only if the state prediction enters B_r with no prior constraint violation. The two operations, prediction and constraint violation check, can be implemented efficiently. A simple method to determine the radius r of a ball $B_r \subseteq \mathbb{X}_K$ is available but omitted for brevity.

The horizon increase made in Step 3) is worthy of a brief discussion, as the heuristic used in increasing N is closely linked to the number of cycles of Algorithm 1 to termination. The simplest approach is to increment N at each iteration of the algorithm, as suggested in [16]. This leads to identification of $N_\infty(x)$ but is likely to result in rather slow implementation of the control scheme. As $\pi^*(x) = \pi_N(x)$ for all $N \geq N_\infty(x)$, however, only an upper bound on $N_\infty(x)$ is required and values of N may therefore be skipped in an effort to obtain a quicker solution.

An interesting approach is to increase N geometrically, starting from a nonzero initial horizon. Assume that the current state is x , that we start with $N = N_0 \geq 1$, and that we multiply the horizon by an integer $k \geq 2$ at each cycle of Algorithm 1. Then termination occurs with $N = k^{\lambda(x)}N_0$ after $\lambda(x) + 1$ cycles, where λ is an integer-valued function of x . Let $c(N)$ be the computational price of solving Problem 2 with horizon N , and further assume that $c(N) \sim O(N)$, so that there exists a finite real $\mu > 0$ such that $c(N) \leq \mu N$, which is a reasonable assumption with infeasible interior point methods [18]. The computational cost of Algorithm 1 is then

$$\begin{aligned} C(x) &= c(N_0) + c(kN_0) + \dots + c[k^{\lambda(x)}N_0] \\ &\leq \mu N_0 [1 + k + \dots + k^{\lambda(x)}] \\ &= \mu N_0 \frac{k^{\lambda(x)+1} - 1}{k - 1}. \end{aligned} \quad (17)$$

If $N_0 \geq N_\infty(x)$, $\lambda(x) = 0$, and $C(x) \leq \mu N_0$. If $N_0 < N_\infty(x)$, we get termination with $N \leq kN_\infty(x) - k$ so that $k^{\lambda(x)}N_0 \leq kN_\infty(x) - k$. It then follows that

$$C(x) \leq \frac{k^2}{k-1} \mu N_\infty(x) - \frac{1}{k-1} \mu (k^2 + N_0). \quad (18)$$

For any choice of N_0 , therefore, we have

$$C(x) \leq \mu \max \left[N_0, \frac{k^2}{k-1} N_\infty(x) - \frac{1}{k-1} (k^2 + N_0) \right]. \quad (19)$$

Consequently, we find that $C(x) \sim O[N_\infty(x)]$. Also, it appears wise to set k to its minimal value 2, leading to

$$C(x) \leq \mu \max [N_0, 4N_\infty(x) - 4 - N_0]. \quad (20)$$

Now, assume $N_\infty(x)$ is known; then calculation of $\pi^*(x)$ requires only the one-shot solution of Problem 2 with $N = N_\infty(x)$, resulting in a minimal computational price $c[N_\infty(x)] \leq \mu N_\infty(x)$. When $N_\infty(x)$ is large (i.e., when computational time is an issue), a good approximation is $c[N_\infty(x)] \approx \mu N_\infty(x)$. In view of (18), we then find that the price of Algorithm 1 is better than approximately $k^2/(k-1)$ times the minimal computational cost required to calculate $\pi^*(x)$.

Another alternative is to use the time-index of the last constraint violation detected at the previous cycle of the algorithm as the new




Fig. 1. Constrained LQR for Van de Vusse reactor.

horizon N . This approach seems to make good control sense, and simulations indicate that it is a reasonable heuristic.

B. A Modification of Algorithm 1

The stability guarantees of Theorems 1 and 2 apply only to the control law given by the exact solution of Problem 1. The stability results are therefore invalid if there is not sufficient time to cycle Algorithm 1 to termination. For cases when time restrictions do not allow termination, Sznaier and Damborg suggest using the last control sequence calculated by the algorithm. They, however, point out that this approach may not lead to stabilizing control; furthermore, it is likely to cause constraint violations.

A modification of Algorithm 1 can, however, be proposed, that leads to guaranteed stability, both nominally and under decaying perturbations, regardless of whether or not the algorithm terminates.

Let MPC_s denote a stabilizing (suboptimal) MPC scheme for the system under consideration. We may then modify *Step 3*), in Algorithm 1, to the following.

Modified Step 3)—Algorithm 1:

Step 3(a): Increase N

Step 3(b): If $N \leq N_{\max}$, go to Step 1).

Step 3(c): Terminate and implement MPC_s. \square

Implementation of Algorithm 1, with the above modification of Step 3), is guaranteed to lead to stabilizing control, even if the algorithm does not terminate in Step 4); of course, when there is sufficient computation time, the algorithm does terminate in Step 4), and this leads to optimal constrained LQR performance.

The stabilizing scheme, MPC_s, may, for instance, be the suboptimal, infinite horizon, stabilizing controller of Rawlings and Muske [11], or a finite-horizon MPC scheme, similar to that discussed in [5]. These control laws provide feasible points for Problem 1 that are suboptimal for any horizon N ; therefore, a by-product of Modified Step 3) is also to give an upper bound for the optimal cost, which may be used to derive an upper bound for N_{∞} .


V. EXAMPLES

A. Van de Vusse Reactor

We consider the isothermal continuous stirred-tank reactor (CSTR) using dilution rate as the manipulated variable. The reactor has constant volume and its dynamics are described by

$$\begin{aligned}\dot{x}_1 &= -k_1 x_1 - k_3 x_1^2 - x_1 u \\ \dot{x}_2 &= k_1 x_1 - k_2 x_2 - x_2 u\end{aligned}\quad (21)$$

which models the Van de Vusse series of reactions



with x_1 and x_2 representing the concentration of A and B and u the dilution (feed) rate [17]. We assume that $k_1 = 50$, $k_2 = 100$, $k_3 = 10$. The control objective is to maintain x_2 at a set-point of one. Consequently, the desired steady states for x_1 and u are 2.5 and 25, respectively [14], [7]. We linearize (21) around this desired steady state and discretize the result with a sampling time of 0.002. This gives

$$A = \begin{bmatrix} 0.95123 & 0 \\ 0.08833 & 0.81873 \end{bmatrix}, \quad B = \begin{bmatrix} -0.0048771 \\ -0.0020429 \end{bmatrix}.$$

We implement the constrained LQR law on this linear system, with an appropriate shift of the origin to account for the nonzero set-point. The tuning we choose is $Q = I$, $R = I$ and the constraint is specified that x_2 should be no greater than 0.12. A simulation is performed that is started from $x_0 = [0.5 \ 0.1]'$ (in the shifted coordinates) and the results are presented in Fig. 1. In this example, we obtain $N_{\infty}(x_0) = 7$ and $\phi^*(x_0) = \phi_7(x_0) = 143.8$.

For comparison, we also present simulation results with two other stabilizing MPC laws. First, we consider an MPC scheme with end-point constraint [5], in which the input parameterization is such that $u = 0$ and $x = 0$ after a finite horizon N . For this control law, if a small horizon N is used, the end-point constraint on the state can become incompatible with the constraint that $x_2 \leq 0.12$. The minimum horizon required for feasibility is, of course, a function of the state. For instance, with $x_0 = [1 \ 0.1]'$, N must be no smaller than eight and with $x_0 = [2 \ 0.1]'$, no smaller than 11.

Returning to the initial condition $x_0 = [0.5 \ 0.1]'$, the minimum horizon required for feasibility is $N = 5$. The open and closed-loop costs of the control strategy obtained with different horizons, from this initial condition are plotted in Fig. 2. (The open-loop cost is the cost associated with the control profile postulated at the first sample, i.e., $\sum_{j=0}^{\infty} x'_{j|0} Q x_{j|0} + u'_{j|0} R u_{j|0}$; the closed-loop cost is the cost associated with the actual controls that are implemented in the receding-horizon implementation, i.e., $\sum_{i=0}^{\infty} x'_i Q x_i + u'_i R u_i$.)

Considering these results, we find for the end-point constrained MPC that we have a difficult problem with the design of N , for which the minimum value depends on the system state, and we also obtain a cost that can be an order of magnitude worse, both in open and closed loop, than with constrained LQR.

We now consider an MPC scheme with infinite costing [11], in which the input parameterization is simply $u = 0$ after a finite horizon N , and there is no end-point state constraint. For this control law, with a horizon $N = 1$, the optimal cost (163.92) is approximately 14% worse than with constrained LQR. With larger N , performance becomes similar to constrained LQR. However, it must be noted that implementation of a stabilizing control law of this type requires enforcement of the state constraint over an infinite horizon, which




Fig. 2. Cost comparison of constrained LQR and MPC with end-point constraint for Van de Vusse reactor.




Fig. 3. Constrained LQR for double integrator, $x_0 = [0.2 \ 0.2]^T$.




Fig. 4. MPC for double integrator, $x_0 = [0.2 \ 0.2]^T$.

leads to solving a finite sequence of quadratic programs just as in the constrained LQR case [6].

We conclude for this example that constrained LQR not only leads to improved performance over alternative MPC formulations, it also releases the user from the design of N . Finally, we point out that the observations made in this example are not specific to the process we use; they would apply to many stable systems.

B. Double Integrator

We consider a double integrator system sampled at a frequency of 10 Hertz, for which

$$A = \begin{bmatrix} 1 & 0 \\ 0.1 & 1 \end{bmatrix}, \quad B = \begin{bmatrix} 0.1 \\ 0.005 \end{bmatrix}.$$

The constraint is specified that the input should not exceed amplitude ten. The constrained LQR law is implemented with $Q = I$, $R = I$.

We consider two different initial conditions, $x_0 = [0.2 \ 0.2]^T$ and $x_0 = [20 \ 20]^T$; a representative simulation is shown in Fig. 3. With

$x_0 = [0.2 \ 0.2]^T$, we get $N_\infty(x_0) = 0$ and $\phi^*(x_0) = 2.23$; with $x_0 = [20 \ 20]^T$, $N_\infty(x_0) = 33$, and $\phi^*(x_0) = 60.056$. We compare with the two stabilizing MPC laws used in Example 1. Because both modes of the system are unstable, the two MPC formulations are equivalent. From initial condition $x_0 = [0.2 \ 0.2]^T$, $N = 4$ provides a feasible problem. The open and closed-loop MPC costs (70, 118) are 30 to 50 times worse than with constrained LQR. Simulation results are displayed in Fig. 4. From initial condition $x_0 = [20 \ 20]^T$, the minimum horizon, after which the state can be forced to zero, is $N = 60$. We implement the control law with this horizon and get performance that compares favorably with constrained LQR; open and closed-loop MPC costs are about 2% larger than the constrained LQR cost. However, the dimension of the optimization problem required to obtain this control law is almost twice as large as that needed in constrained LQR.

To summarize this example we have the following.

- 1) For operation close to the steady state, constrained LQR uses control that is not even close to the constraints. With MPC, a

small N can be used for online efficiency, but that choice leads to large inputs which are unnecessary and undesirable when compared to those obtained with constrained LQR. Moreover, the computational cost of implementing MPC is higher than constrained LQR.

- 2) For operation away from steady state, a larger number of decision variables (almost twice as many) is needed for MPC compared to constrained LQR. Although MPC can be implemented in this case by solving a single quadratic program (QP), the dimension of that QP is much larger than any of the QP's solved in the constrained LQR approach. Again, the computational cost for MPC is higher and the performance is worse.

VI. CONCLUSIONS

In this paper, we further developed the constrained LQR problem proposed by Sznaier and Damborg. In addition to being optimal, the constrained LQR removes the mismatch between open and closed-loop nominal behavior, which is an undesirable side effect of current MPC approaches that are based on finite-input parameterizations. The main practical advantage of the constrained LQR compared to currently available stabilizing MPC approaches, however, is not optimality but simplicity. The constrained LQR removes the tuning parameter N , for which no reliable tuning guidelines are available. Our recommendation is that current industrial practice in which control engineers purposefully use a suboptimal and small value of N in order to decide the performance robustness tradeoff should be reconsidered. The performance robustness tradeoff can be decided, as in the classic unconstrained LQR case, by choosing appropriate Q and R .

We presented the outline of an algorithm that allows practical implementation of the constrained LQR; the computational cost of this algorithm is reasonable compared to the minimal cost required for calculation of the optimal solution. However, examples exist for which this minimal cost is large, which may preclude real-time implementation of the constrained LQR in some situations. It is not clear if these cases arise often enough in practice to warrant serious attention, but the algorithm is easily modified so that termination occurs for such cases automatically with a stabilizing MPC solution, and without user intervention or online tuning.

It has been assumed throughout this paper that the state and input constraints are compatible, i.e., $\mathbb{IP} \neq \emptyset$. Even though the assumption that $h \in \mathbb{R}_+^{n_h}$ and $d \in \mathbb{R}_+^{n_d}$ ensures that steady state at the origin is feasible, the problem of transient infeasibility remains. Relevant contributions have been made in the context of MPC and constraint relaxation techniques have been proposed that do not endanger closed-loop stability. These are readily transportable to constrained LQR.

After this paper was submitted, the authors obtained a recent report of related work by Chmielewski and Manousiouthakis [2]. The main differences are as follows. As in [16], Chmielewski and Manousiouthakis assume the state and input constraint sets are compact, convex polyhedrons. In this paper we do not assume these constraint sets are bounded. In many applications there are no constraints specified on some states or inputs and the constraint regions are therefore unbounded. Because the constrained LQR problems differ, the algorithms for computing solutions also differ. The compact state constraint set enables Chmielewski and Manousiouthakis to compute an upper bound on N_∞ and solve a single QP. In this paper, because x_0 is not in a compact set, N is increased in a series of QP's until termination. For cases with compact constraint regions, further research is required to make a quantitative comparison between the efficiencies of these two approaches.

ACKNOWLEDGMENT

The authors would like to thank Dr. S. J. Wright for discussion of the computational cost of the algorithm with infeasible interior point methods.

REFERENCES

- [1] R. R. Bitmead, M. Gevers, and V. Wertz, *Adaptive Optimal Control, The Thinking Man's GPC*. Englewood Cliffs, NJ: Prentice-Hall, 1990.
- [2] D. Chmielewski and V. Manousiouthakis, "On constrained infinite-time linear quadratic optimal control," *Syst. Contr. Lett.*, vol. 29, pp. 121–129, 1996.
- [3] C. R. Cutler and B. L. Ramaker, "Dynamic matrix control—A computer control algorithm," in *Proc. Joint Automatic Control Conf.*, 1980.
- [4] R. E. Kalman, "Contributions to the theory of optimal control," *Bull. Soc. Math. Mex.*, vol. 5, pp. 102–119, 1960.
- [5] S. S. Keerthi and E. G. Gilbert, "Optimal infinite-horizon feedback laws for a general class of constrained discrete-time systems: Stability and moving-horizon approximations," *J. Optim. Theory Appl.*, vol. 57, no. 2, pp. 265–293, May 1988.
- [6] E. S. Meadows, K. R. Muske, and J. B. Rawlings, "Implementable model predictive control in the state space," in *Proc. American Control Conf.*, 1995, pp. 3699–3703.
- [7] E. S. Meadows and J. B. Rawlings, "Model predictive control," in *Nonlinear Process Control*, M. A. Henson and D. E. Seborg, Eds. Englewood Cliffs, NJ: Prentice Hall, 1997, pp. 233–310.
- [8] H. Michalska and D. Q. Mayne, "Robust receding horizon control of constrained nonlinear systems," *IEEE Trans. Automat. Contr.*, vol. 38, pp. 1623–1633, Nov. 1993.
- [9] K. R. Muske, "Linear model predictive control of chemical processes," Ph.D. dissertation, Univ. Texas, Austin, 1995.
- [10] J. B. Rawlings, E. S. Meadows, and K. R. Muske, "Nonlinear model predictive control: A tutorial and survey," in *ADCHEM'94 Proc.*, Kyoto, Japan.
- [11] J. B. Rawlings and K. R. Muske, "Stability of constrained receding horizon control," *IEEE Trans. Automat. Contr.*, vol. 38, pp. 1512–1516, Oct. 1993.
- [12] J. Richalet, A. Rault, J. L. Testud, and J. Papon, "Model predictive heuristic control: Applications to industrial processes," *Automatica*, vol. 14, pp. 413–428, 1978.
- [13] P. O. Scokaert, J. B. Rawlings, and E. S. Meadows, "Discrete-time stability with perturbations: Application to model predictive control," *Automatica*, vol. 33, no. 3, pp. 463–470, 1997.
- [14] P. B. Sistu and B. W. Bequette, "Model predictive control of processes with input multiplicities," *Chem. Eng. Sci.*, vol. 6, no. 6, pp. 921–936, 1995.
- [15] M. Snazier, "Suboptimal feedback control of constrained linear systems," Ph.D. dissertation, Univ. Washington, 1989.
- [16] M. Sznaier and M. J. Damborg, "Suboptimal control of linear systems with state and control inequality constraints," in *Proc. 26th Conf. Decision and Control*, 1987, pp. 761–762.
- [17] J. G. Van de Vusse, "Plug-flow type reactor versus tank reactor," *Chem. Eng. Sci.*, vol. 19, p. 964, 1964.
- [18] S. J. Wright, "Applying new optimization algorithms to model predictive control," in *Fifth Int. Conf. Chemical Process Control*, Lake Tahoe, CA, Jan. 1996.