Statement 2: Under the assumptions of this theorem for any $0 \le t_0 \le t$

$$\Phi_{\tilde{u}}(t) - \Phi_{\tilde{u}}(t_0) \ge \Phi_{\tilde{u}*}(t) - \Phi_{\tilde{u}*}(t_0)$$
$$= -\frac{1}{h}\Delta\vartheta^T(t)L^{-1}\Delta\vartheta(t) + o_\omega(h)$$

where

$$\Phi_{\tilde{u}}(t) := \int_0^t (2\tilde{u}^T(\tau)d\vartheta(\tau) + \tilde{u}^T(\tau)L\tilde{u}(\tau)d\tau) \qquad (20)$$
$$\Delta\vartheta(t) := B_{u0}^T P[B_{u0}(u_{\mathrm{comp}}(t) - x^*(t))h$$
$$+ (K + A_0 C_0^+)\Delta y(t)], \quad \Delta y(t) = y(t) - y(t-h)$$

and this minimum is reachable for

$$\tilde{u}(t)\,dt = \tilde{u}^*(t)\,dt := -L^{-1}B_{u0}^T P[B_{u0}(u_{\mathrm{comp}}(t) + x^*(t))\,dt$$
$$+ (K + A_0 C_0^+)\,dy(t)].$$

As a result, we have: $d\Phi_{\tilde{u}*}(t) \le 0$ (in symbolic form).

Proof of Statement 2: Using the Euler–Maruyama's formula [3], [9] we obtain the following relation: $(h := t - t_0 \to 0)$: $\Phi_{\tilde{u}}(t) - \Phi_{\tilde{u}}(t_0) = 2\tilde{u}^T(t)\Delta\vartheta(t) + \tilde{u}^T(t)L\tilde{u}(t)h + o_\omega(h)$.

Minimizing then the right side for each fixed t, we derive

$$\min_{\tilde{u}(t)}[\Phi_{\tilde{u}}(t) - \Phi_{\tilde{u}}(t_0)] = \Phi_{\tilde{u}*}(t) - \Phi_{\tilde{u}*}(t_0)$$
$$= -\frac{1}{h}\Delta\vartheta^T(t)L^{-1}\Delta\vartheta(t) + o_\omega(h) \le o_\omega(h). \qquad (21)$$

Hence, taking into account the definition (20), we have $d\Phi_{\tilde{u}*}(t) \le 0$. □

Rewriting (18) in differential form and taking into account that P is the solution of the Riccati equation (see Assumption 4 of this theorem) we finally derive

$$dV(e(t)) \overset{\mathrm{a.s.}}{\le} [\varphi(t) + I(t)]\,dt + S^T(t)\,dw(t) - e^T(t)Q_e e(t)\,dt.$$

According to the Kronecker lemma, the Large Number law for martingale, and Skorohod lemma (see [2] and [4]), we obtain the result of the theorem. Details of this proof can be found in [13].

REFERENCES

[1] S. Boyd, L. Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory.* Philadelphia, PA: SIAM, 1994.

[2] H. Cramer and M. Lidbetter, *Stationary Random Processes.* Prinstone, 1969.

[3] T. Gard, *Introduction to Stochastic Differential Equations.* New York: Marcel Dekker, 1988.

[4] I. I. Gihman and A. V. Skorohod, *Stochastic Differential Equations.* New York: Springer-Verlag, 1972.

[5] M. Grimble, "Three and half DOF polynomial solution of the feedforward H_2/H_∞ control problem," in *Proc. 34rd Conf. Decision and Control*, New Orleans, LA, Dec. 1995, pp. 4145–4150.

[6] K. Gu, "H_∞ control of systems under norm bounded uncertainties in all system matrices," *IEEE Trans. Automat. Contr.*, vol. 39, June 1994.

[7] R. Lipster and A. Shiryayev, *Statistics of Random Processes*, vols. 1, 2. New York: Springer-Verlag, 1977.

[8] R. Mangoubi, "Robust estimation and failure detection for linear systems," Ph.D. dissertation, Dept. Aeronautic and Astronautic, MIT 1995.

[9] G. Maruyama, "Continuous markov processes and stochastic equations," *Rend. Circ. Mat. Palermo.*, vol. 4, pp. 48–90.

[10] Y. Nesterov and A. Nemirovsky, "Interior-point polynomial methods in convex programming," *Studies in Appl. Math. SIAM*, vol. 13, 1994.

[11] B. T. Polyak, "Introduction to optimization," *Optimization Software, Publication Division.* New York, 1987.

[12] A. Poznyak and M. Taksar, "Robust state feedback stochastic control for linear systems with time-varying parameters," in *Proc. 33rd Conf. Decision and Control*, Lake Buena Vista, FL, Dec. 1994, pp. 1161–1162.

[13] ——, "Robust control of linear stochastic systems with fully observable state," *Applicationes Mathematicae*, vol. 24, no. 1, pp. 35–46, 1996.

[14] M. I. Taksar, A. S. Poznyak, and A. Iparraguirre, "Robust output feedback control for linear stochastic systems in continuous time with time-varying parameters," in *Proc. 34rd Conf. Decision and Control*, New Orleans, LA, Dec. 1995, pp. 2179–2184.

[15] J. C. Willems, "Least squares stationary optimal control and algebraic Riccati equation," *IEEE Trans. Automat. Contr.*, vol. AC-16, no. 6, pp. 621–634, 1971.

Min–Max Feedback Model Predictive Control for Constrained Linear Systems

P. O. M. Scokaert and D. Q. Mayne

Abstract—Min–max feedback formulations of model predictive control are discussed, both in the fixed and variable horizon contexts. The control schemes the authors discuss introduce, in the control optimization, the notion that feedback is present in the receding-horizon implementation of the control. This leads to improved performance, compared to standard model predictive control, and resolves the feasibility difficulties that arise with the min–max techniques that are documented in the literature. The stabilizing properties of the methods are discussed as well as some practical implementation details.

Index Terms—Feedback, min–max optimization, model predictive control.

I. INTRODUCTION

Model predictive control (MPC) is a control methodology that is becoming mature. The basic aspects of the method are now well understood, and stabilizing formulations of the control law are documented in the literature, both for linear and nonlinear processes [1], [2].

The MPC strategy optimizes an open-loop control sequence, at each sample, to minimize a nominal cost function, subject to some state and input constraints. The optimization is usually based on the assumption that the process model is exact and that future disturbances are constant. Because the control law ignores the effects of possible future changes in disturbance and model mismatch, closed-loop performance can be poor [3], with likely violations of the constraints, when disturbances or model mismatch are present.

Some formulations of MPC have been proposed that address these issues [4], [5]. These methods rely on a min–max optimization of predicted performance. However, these formulations optimize a single control profile over all possible disturbance (or model mismatch) sequences and therefore do not include the notion that feedback is present in the receding-horizon implementation of the control. When

unknown disturbances affect the process, the result is often that the control law is inapplicable, due to infeasibility.

In this paper, we consider a min–max feedback MPC approach, in which the model is assumed to be exact, and the effects of possible future disturbances are taken into account. The formulation we propose does introduce the notion of feedback in the control optimization that is performed at each sample.

II. FEEDBACK MIN–MAX MPC

We consider linear time-invariant discrete-time systems described by

$$x_{t+1} = Ax_t + Bu_t + w_t \tag{1}$$

where $x_t \in \mathbf{R}^n$, $u_t \in \mathbf{R}^m$ are the state and input vectors, at time t; A and B are the state transition and input distribution matrices; and we assume that (A, B) is controllable. The vector $w_t \in \mathbf{W} \subset \mathbf{R}^n$ is an unknown, but bounded, deterministic disturbance; $\mathbf{W}$ is compact, convex, and contains an open neighborhood of the origin.

Because the system cannot be controlled to the origin, due to w_t, the standard MPC law must be modified. An obvious possibility is to ignore the disturbance. However, the resulting performance can be poor [3]. To overcome this deficiency, we may consider the use of min–max MPC formulations, similar to those proposed in [4] and [5]. However, this raises other difficulties. On one hand, the usual stability constraint $x_{t+N} = 0$ cannot be included in the formulation, since the system cannot be steered to the origin. Also, the worst case scenarios corresponding to a nominal control can be considerably worse than actually achieved by feedback control. This is highly likely to lead to infeasibility, when state constraints are used.

In this paper, we consider an MPC variant that is also based on a min–max optimization of the control, but in which we include the notion that feedback is present. The result is a control law that is more computationally intensive but which leads to more reliable and consistent performance.

A. The Control Law

As the state cannot be steered to the origin, the control objective is to regulate the state of the system to a predetermined robust control invariant set. A set of state and input constraints must also be satisfied. We denote the robust control invariant set by $\mathbf{X}_0$ and the constraints are given by

$$x_t \in \mathbf{X}, \qquad u_t \in \mathbf{U} \tag{2}$$

where $\mathbf{X}_0$, $\mathbf{X}$, and $\mathbf{U}$ are compact, convex sets, each containing an open neighborhood of the origin, and $\mathbf{X}_0 \subseteq \mathbf{X} \subseteq \mathbf{R}^n$, $\mathbf{U} \subseteq \mathbf{R}^m$.

We consider a dual-mode control law, for which we design both an "inner" and an "outer" controller. The inner controller operates when the state is in the robust control invariant set, $\mathbf{X}_0$, and its role is to keep the state in this set, despite the disturbance. The outer controller operates when the state is outside the invariant set and steers the system state to the invariant set.

The inner controller we use is linear, of the form $u = -Kx$, and is such that $(A - BK)^s = 0$, for some finite integer s. This property of the inner controller is important in the construction of the control robust invariant set, as seen in Section II-B.

For the outer controller, we propose two feedback min–max MPC schemes, which form the focus of this paper. First, we consider a fixed horizon formulation; then we present, in Section III, a variable horizon variant of the method.

We start with the fixed horizon formulation of the control law. At time t, let $\{w_{j|t}^\ell\}$ denote the possible realizations of the disturbance; $\ell \in \mathcal{L}$ indexes the realizations. Further, let $\{u_{j|t}^\ell\}$ denote a control

sequence associated with the ℓth such realization, and let $x_{t+j|t}^\ell$ represent the solution of the model equation

$$x_{j+1|t}^\ell = Ax_{j|t}^\ell + Bu_{j|t}^\ell + w_{j|t}^\ell, \qquad \ell \in \mathcal{L} \tag{3}$$

with $x_{t|t}^\ell = x_t$, for all $\ell \in \mathcal{L}$. In the fixed horizon case, we consider the min–max optimization

$$\min_{\{u_{j|t}^\ell\}} \max_{\ell \in \mathcal{L}} \sum_{j=0}^{N-1} L\left(x_{t+j|t}^\ell, u_{t+j|t}^\ell\right)$$

$$\begin{aligned}
\text{s.t.} \quad & x_{j|t}^\ell \in \mathbf{X}, & j > t, \quad \forall \ell \in \mathcal{L}, \\
& u_{j|t}^\ell \in \mathbf{U}, & j \geq t, \quad \forall \ell \in \mathcal{L}, \\
& x_{t+N|t}^\ell \in \mathbf{X}_0, & \forall \ell \in \mathcal{L}, \\
& x_{j|t}^{\ell_1} = x_{j|t}^{\ell_2} \Rightarrow u_{j|t}^{\ell_1} = u_{j|t}^{\ell_2}, & j \geq t, \quad \forall \ell_1, \ell_2 \in \mathcal{L}
\end{aligned} \tag{4}$$

where N is the control horizon, and the stage cost, $L: \mathbf{R}^n \times \mathbf{R}^m \to \mathbf{R}$, is a positive semidefinite function. The outer controller is obtained by receding-horizon implementation of the solution to this optimization.

The terminal constraint $x_{t+N|t}^\ell \in \mathbf{X}_0$ is the stability constraint for this formulation and is necessary in order to guarantee stability. The notion of feedback is introduced into the control optimization by allowing a different control sequence for each realization of the disturbance. Combined with this, however, is the "causality constraint," $x_{j|t}^{\ell_1} = x_{j|t}^{\ell_2} \Rightarrow u_{j|t}^{\ell_1} = u_{j|t}^{\ell_2}$, which restricts the freedom on the control sequences considered by enforcing a single control action for each state. Thus, the control predicted for time j depends only on the state prediction for time j, not on the path taken to reach that prediction.

Note that in the min–max formulation of [4] and [5] a single control sequence is optimized over all disturbance realizations. Thus, the notion that more measurements become available as time progresses is not included in the formulation; this is the cause of the feasibility difficulties that are likely to arise with that formulation.

B. Construction of the Robust Control Invariant Set

A set $\mathbf{X} \subset \mathbf{R}^n$ is said to be robust control invariant if system (1), with the predetermined control law $u = -Kx$, satisfies the state and input constraints in $\mathbf{X}$ and if $(A - BK)\mathbf{X} + W \subset \mathbf{X}$. Thus, if the initial state, x, lies in robust a control invariant set, the control law, $u = -Kx$, keeps all subsequent states in this set, despite the disturbance.

The first step in the construction of the robust control invariant set, $\mathbf{X}_0$, is to define the inner controller. We may use any linear controller, $u = -Kx$, that steers the system state to the origin in finite time, when the disturbance is absent. Such a controller always exists because (A, B) is assumed controllable. For all appropriate designs of K, we thus get an integer $s \geq n$ such that $(A - BK)^s = 0$.

The inner controller is used as soon as the state enters $\mathbf{X}_0$. For all initial states $x_t \in \mathbf{X}_0$, we therefore get $x_{t+1} = Fx_t + w_t$, where $F = A - BK$. By design of K, we also have $F^s x = 0$, $F^s w = 0$, for all $x, w \in \mathbf{R}^n$. If $x_{t+j} \in \mathbf{X}_0$ for all $j \geq 0$, it follows that, with $j \geq s$, $x_{t+j} = w_{t+j-1} + Fw_{t+j-2} + \cdots + F^{s-1}w_{t+j-s}$. In view of this, we define the set

$$\mathbf{X}_0 = \mathbf{W} + F\mathbf{W} + \cdots + F^{s-1}\mathbf{W}. \tag{5}$$

This set is robust control invariant if $\mathbf{X}_0 \subseteq \mathbf{X}$ and $-K\mathbf{X}_0 \subseteq \mathbf{U}$.

C. The Control Algorithm

The fixed horizon feedback min–max MPC law we propose is summarized as follows.

Algorithm 1—Fixed Horizon Feedback Min–Max MPC:

Data: x_t.

Algorithm: If $x_t \in \mathbf{X}_0$, set $u_t = -Kx_t$. Otherwise, find the solution of (4) and set u_t to the first control in the optimal sequence calculated.

D. Properties

Certain properties follow immediately from the formulation of the control optimization. The worst case performance obtained with this control law, i.e., the performance obtained with the disturbance realization that is most "upsetting" to the system, has a cost no greater than the optimal cost for a min–max approach that does not incorporate the notion of feedback. This results from the increased number of degrees of freedom in the control optimization.

Under certain assumptions on the cost, the fixed horizon feedback min–max control law can also be shown to be stabilizing.

Assumption 1: The function L is convex and such that

$$
\begin{aligned}
L(x,u) &= 0, & \forall x \in \mathbf{X}_0, u = -Kx \\
L(x,u) &\geq \alpha(d(x,\mathbf{X}_0)), & \forall x \notin \mathbf{X}_0, \forall u
\end{aligned} \tag{6}
$$

where α is a K-function and $d(x,\mathbf{X}_0)$ denotes the distance of a point x to the set $\mathbf{X}_0$.

Theorem 1: Suppose Assumption 1 holds. Then, the feedback min–max MPC law given by Algorithm 1 drives the state x_t asymptotically to the robust control invariant set $\mathbf{X}_0$.

Proof: At time t, state x_t, let $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}$, represent the possible disturbance realizations and let $\{\bar{u}_{j|t}^\ell\}$ denote the corresponding optimal control sequences. Further, let $\{x_{j|t}^\ell\}$ and ϕ_t^ℓ denote the state trajectories and costs associated with each of the optimal control/disturbance realizations. The optimal cost, at time t, is $\bar{\phi}_t = \max_{\ell \in \mathcal{L}} \phi_t^\ell$. At time t, the first of the optimal controls is applied and the disturbance takes a certain value w_t. Let $\mathcal{L}_1$ denote the set of indexes such that $w_{t|t}^\ell = w_t$ for all $\ell \in \mathcal{L}_1$, and $w_{t|t}^\ell \neq w_t$ for all $\ell \notin \mathcal{L}_1$. At time $t+1$, the state has moved along a trajectory that coincides with the predicted state trajectories indexed by $\ell \in \mathcal{L}_1$. The control sequences $[\bar{u}_{t+1|t}^\ell, \bar{u}_{t+2|t}^\ell, \cdots, \bar{u}_{t+N-1|t}^\ell, -Kx_{t+N|t}^\ell]$, $\ell \in \mathcal{L}_1$, then satisfy the constraints and yield costs $\phi_t^\ell - L(x_t,u_t) + L(x_{t+N|t}^\ell, -Kx_{t+N|t}^\ell) = \phi_t^\ell - L(x_t,u_t)$, since $x_{t+N|t}^\ell \in \mathbf{X}_0$ for all $\ell \in \mathcal{L}$. The optimal cost at time $t+1$ is no larger than the largest of these costs; denoting by $\bar{\phi}_{t+1}$ the optimal cost at time $t+1$, we therefore get $\bar{\phi}_{t+1} \leq \max_{\ell \in \mathcal{L}_1} \phi_t^\ell - L(x_t,u_t)$, and consequently

$$
\bar{\phi}_{t+1} \leq \bar{\phi}_t - L(x_t, u_t) \tag{7}
$$

since $\max_{\ell \in \mathcal{L}_1} \bar{\phi}_t^\ell \leq \max_{\ell \in \mathcal{L}} \bar{\phi}_t^\ell = \bar{\phi}_t$. The cost is therefore monotonically nonincreasing. As it is bounded below by zero, it must consequently converge to a constant value, so that $\bar{\phi}_t - \bar{\phi}_{t+1} \to 0$, as $t \to \infty$. From (7), we have $L(x_t,u_t) \leq \bar{\phi}_t - \bar{\phi}_{t+1}$, and it follows that $L(x_t,u_t) \to 0$, as $t \to \infty$. In view of Assumption 1, we conclude that $d(x_t,\mathbf{X}_0) \to 0$, as $t \to \infty$, i.e., the state converges to $\mathbf{X}_0$ asymptotically.   □

Remark 1: Note that, if the state enters $\mathbf{X}_0$ at any time, it is subsequently kept inside this set by the inner controller $u = -Kx$.

Remark 2: Note also that if, in addition to Assumption 1, we have $L(x,u) \geq \alpha(\|x\|)$, for all $x \notin \mathbf{X}_0$, $u \in \mathbf{R}^m$, then the state can be shown to enter $\mathbf{X}_0$ in finite time.

The min–max optimization for the outer controller associates a separate control sequence for the state trajectories that can result from every possible disturbance realization. In general, prior knowledge of the disturbance is limited to a compact convex set $\mathbf{W}$, within which w_t may be assumed to lie. This means an infinite number of disturbance realizations must be considered, leading to an optimization of infinite dimension. Practical implementation of the controller then appears possible only in an approximate sense, through quantization of the disturbance; moreover, this type of approximation invalidates the proof of Theorem 1. However, due to linearity of the process and convexity of the constraints and cost, this problem can be resolved. Indeed, we show below that if $\mathbf{W}$ is a polytope in $\mathbf{R}^n$, considering, in the control optimization, only the "extreme" disturbance realizations, which take values at the vertices of $\mathbf{W}$, it leads to a stabilizing control law that satisfies the constraints for all disturbance realizations that lie inside the convex hull of the realizations considered in the optimization.

Remark 3: An "extreme disturbance realization" is a sequence which takes values at the vertices of the convex set $\mathbf{W}$. Therefore, every disturbance realization that takes values in $\mathbf{W}$ lies in the convex hull of the extreme disturbance realizations.   □

Again, at time t, let $\{w_{j|t}^\ell\}$ denote the possible realizations of the disturbance, with $\ell \in \mathcal{L}$. Further, let $\mathbf{W}$ be a polytope in $\mathbf{R}^n$ and let $\mathcal{L}_v$ denote the set of indexes ℓ, such that $\{w_{j|t}^\ell\}$ takes values only on the vertices of $\mathbf{W}$. We now consider the min–max control optimization

$$
\min_{\{u_{j|t}^\ell\}} \max_{\ell \in \mathcal{L}_v} \sum_{j=0}^{N-1} L\left(x_{t+j|t}^\ell, u_{t+j|t}^\ell\right)
$$

$$
\begin{aligned}
\text{s.t.} \quad & x_{j|t}^\ell \in \mathbf{X}, & j > t, \quad \forall \ell \in \mathcal{L}_v, \\
& u_{j|t}^\ell \in \mathbf{U}, & j \geq t, \quad \forall \ell \in \mathcal{L}_v, \\
& x_{t+N|t}^\ell \in \mathbf{X}_0, & \forall \ell \in \mathcal{L}_v, \\
& x_{j|t}^{\ell_1} = x_{j|t}^{\ell_2} \Rightarrow u_{j|t}^{\ell_1} = u_{j|t}^{\ell_2}, & j \geq t, \quad \forall \ell_1, \ell_2 \in \mathcal{L}_v
\end{aligned} \tag{8}
$$

and we obtain the outer control by receding-horizon implementation of the solution to this optimization. We note immediately that this optimization has finite dimension, because $\mathcal{L}_v$ contains only a finite number of indexes; this number depends on the number of vertices of $\mathbf{W}$ and the horizon N. We show that the resulting min–max control law satisfies the state and input constraints and is stabilizing for all disturbance realizations $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}$.

Theorem 2: Suppose Assumption 1 holds, and let $\mathbf{W}$ be a polytope in $\mathbf{R}^n$, with $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}_v$, denoting the extreme disturbance realizations, which take values at the vertices of $\mathbf{W}$. Then, the feedback min–max law given by Algorithm 1, with the optimization of (4) replaced by (8), drives the state x_t asymptotically to the robust control invariant set $\mathbf{X}_0$.

Proof: At time t, state x_t, let $\{\bar{u}_{j|t}^\ell\}$, $\ell \in \mathcal{L}_v$ denote the optimal control sequences corresponding to the disturbance realizations $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}_v$. Further, let $\{x_{j|t}^\ell\}$ denote the state trajectories associated with each of these optimal control/disturbance realizations. Finally, let $\bar{u}_t$ denote the first control in the optimal control sequences. At time t, state x_t, the optimal control $\bar{u}_t$ is applied and the disturbance takes a certain value $w_t \in \mathbf{W}$, driving the state to $x_{t+1} = Ax_t + B\bar{u}_t + w_t$. Due to linearity of the process, x_{t+1} lies in $\text{co}\{x_{t+1|t}^\ell \mid \ell \in \mathcal{L}_v\}$, where $x_{t+1|t}^\ell = Ax_t + B\bar{u}_t + w_{t|t}^\ell$, and $\text{co}\{\cdot\}$ denotes the convex hull of the points in a set $\{\cdot\}$. Therefore, x_{t+1} may be expressed as

$$
x_{t+1} = \sum_{\ell \in \mathcal{L}_v} \mu_\ell x_{t+1|t}^\ell \tag{9}
$$

where μ_ℓ are appropriate scalar weights. At time $t+1$, state x_{t+1}, consider the control sequence defined as

$$
\left[\sum_{\ell \in \mathcal{L}_v} \mu_\ell \bar{u}_{t+1|t}^\ell, \cdots, \sum_{\ell \in \mathcal{L}_v} \mu_\ell \bar{u}_{t+N-1|t}^\ell, \sum_{\ell \in \mathcal{L}_v} -\mu_\ell Kx_{t+N|t}^\ell\right]. \tag{10}
$$

Under this control strategy, the state and input predictions evolve in the convex hulls of the predictions made at time t so that

$$
\begin{aligned}
x_{j|t+1}^\ell &\in \text{co}\{x_{j|t}^\ell \mid \ell \in \mathcal{L}_v\}, \\
u_{j|t+1}^\ell &\in \text{co}\{\bar{u}_{j|t}^\ell \mid \ell \in \mathcal{L}_v\},
\end{aligned} \quad \forall \ell \in \mathcal{L}. \tag{11}
$$

Fig. 1.   Example: nominal MPC, $w_t = -1/t$.

Since $x_{j|t}^\ell \in \mathbf{X}$ and $\bar{u}_{j|t}^\ell \in \mathbf{U}$ for all $\ell \in \mathcal{L}_v, j$, it follows that, under this control sequence, $x_{j|t+1}^\ell \in \mathbf{X}$ and $u_{j|t+1}^\ell \in \mathbf{U}$ for all $\ell \in \mathcal{L}, j$. Also, because $x_{t+N-1|t}^\ell \in \mathbf{X}_0$ for all ℓ and $\mathbf{X}_0$ is robust control invariant under the control law $u = -Kx$, it follows that $x_{t+N|t+1}^\ell \in \mathbf{X}_0$ for all $\ell \in \mathcal{L}$. So we find that the control sequence in (10) leads to satisfaction of the state, input, and stability constraints, at time $t + 1$. Furthermore, it yields a cost that is no larger than the largest cost at time t, because L is convex, making the cost a convex function of the state and input predictions.

As the control sequence in (10), which may be suboptimal, satisfies the constraints and yields a cost that is no larger than the optimal cost at time t, we conclude that the optimal cost, at time $t + 1$, is no larger the optimal cost at time t. Thus, the cost is monotonically nonincreasing with time, and using the same arguments as in the proof for Theorem 1, we conclude that the state converges asymptotically to the set $\mathbf{X}_0$. $\square$

E. Computational Issues

The control optimization may be stated as

$$\min_{\mathbf{u}} \max_{\ell \in \mathcal{L}_v} \{ \phi_t^\ell \mid \mathbf{u} \in \mathbf{C} \} \qquad (12)$$

where

$$\mathbf{u} = (\mathbf{u}^1, \mathbf{u}^2, \cdots) \qquad (13)$$

with $\mathbf{u}^\ell = (u_{t|t}^\ell, u_{t+1|t}^\ell, \cdots, u_{t+N-1|t}^\ell)$; the set $\mathbf{C}$ is defined by the stability, state, and input constraints.

This is a linearly constrained convex min–max optimization, which can be solved using standard optimization techniques. If $\mathbf{W}$ has p vertices, the number of controls that are computed is $m(1 + p + p^2 + \cdots + p^{N-1})$. This is the dimension of the min–max optimization.

An alternative formulation of the control optimization is

$$\min_{\mathbf{u},v} \{ v \mid \phi_t^\ell \le v, \text{ for all } \ell \in \mathcal{L}_v, \mathbf{u} \in \mathbf{C} \} \qquad (14)$$

which is a standard nonlinear program, of dimension $m(1 + p + p^2 + \cdots + p^{N-1}) + 1$, the extra dimension being due to the addition of v as a degree of freedom. We note that although the constraints that define $\mathbf{C}$ are linear, the constraints $\phi_t^\ell \le v$ are not linear inequalities. However, the problem is convex and can therefore be solved by available algorithms.

Finally, we note that when the process model is exact, and the disturbance remains in $\mathbf{W}$, open-loop implementation of the optimal control sequence computed at any time satisfies the constraints and drives the state to the invariant set, $\mathbf{X}_0$, in finite time. Further reduc-

tions in cost may be obtained by receding-horizon implementation of the outer controller; nevertheless reoptimization is not necessary, if computation is at a premium.

F. Illustrative Example

To illustrate the points we make, we consider a very simple example. The system is described as

$$x_{k+1} = x_k + u_k + w_k \qquad (15)$$

with $w_- \le w_k \le w_+$, where w_- and w_+ are known bounds on the disturbance, w_k. The constraints are given by

$$\begin{aligned} x_{j+1|t} &\in \mathbf{X} = \{ x \in \mathbf{R} : -1.2 \le x \le 2 \}, \\ u_{j|t} &\in \mathbf{U} = \mathbf{R}, \end{aligned} \qquad \forall j \ge t. \qquad (16)$$

For this example, we use $-w_- = w_+ = 1$. The inner controller is $u = -x$, and the robust control invariant set is $\mathbf{X}_0 = \{ x \in \mathbf{R}^n : |x| \le 1 \}$. We use a horizon $N = 3$ and the stage cost is $L(x, u) = 0$, if $x \in \mathbf{X}_0$, and $L(x, u) = x^2 + u^2$, otherwise.

Nominal MPC: In nominal MPC, a single control sequence, $[u_{t|t}, u_{t+1|t}, u_{t+2|t}]$, is optimized at each time t, and the effects of the disturbances, w_0, w_1, and w_2, are ignored in the control optimization.

The optimal control ensures that the state constraints are satisfied, only if $w_0 = w_1 = w_2 = 0$. Therefore, we have $x_{t+1|t} \in \mathbf{X}$. However, the actual disturbance may not be zero, and the best approximation we have of x_{t+1} is: $x_{t+1} \in [x_{t+1|t} - w_- \quad x_{t+1|t} + w_+]$. Consequently, we may get violations of the state constraint in closed-loop, even though the controller predicts satisfaction of the constraints at all times.

Also, because x_{t+1} may not coincide with $x_{t+1|t}$, the control sequence $[u_{t+1|t}, u_{t+2|t}, -Kx_{t+3|t}]$ may not satisfy the state constraints at time $t + 1$. Consequently, the stabilizing properties of traditional MPC are lost.

To illustrate this, we use an MPC scheme, in which the control is given by minimization of the nominal objective, subject to the nominal constraints, if $x \notin \mathbf{X}_0$, and by $u = -x$ otherwise. (The nominal objective and constraints are based on the predictions that result by assuming that $w_t = 0$ for all t.)

In Fig. 1, we show the results obtained with $x_0 = 2$ and $w_t = -1/t$, $t \ge 1$; in Fig. 2, the results obtained with $x_0 = 2$ and $w_t = -\cos(t/5)$, $t \ge 0$ are presented. For this example, we find, in both cases, that the state is brought to the robust control invariant set $\mathbf{X}_0$ in finite time. However, constraint violations are experienced during the transient to $\mathbf{X}_0$.

Fig. 2.   Example: nominal MPC, $w_t = -\cos(t/5)$.

Traditional Min–Max MPC: In traditional min–max MPC, a single control sequence, $[u_{t|t}, u_{t+1|t}, u_{t+2|t}]$, is optimized over all possible disturbance profiles, at each time t. We therefore get the prediction

$$x_{t+3|t} = x_t + u_{t|t} + u_{t+1|t} + u_{t+2|t} + w_0 + w_1 + w_2 \qquad (17)$$

which yields $x_{t+3|t} = x_t + u_0 + u_1 + u_2 - 3$, if $w_0 = w_1 = w_2 = -1$, and $x_{t+3|t} = x_t + u_0 + u_1 + u_2 + 3$, if $w_0 = w_1 = w_2 = 1$.

From this, we find that it is impossible to satisfy the state constraint, $x_{j|t} \in \mathbf{X}$, for all disturbance realizations. The control optimization is therefore infeasible and the control law fails.

Even if the state constraint is removed, it is impossible to satisfy the stability constraint, $x_{t+N|t} \in \mathbf{X}_0$, for all disturbance realizations. Consequently, the control law still fails, due to infeasibility. Moreover, attempts to continue control in spite of this problem result in a loss of the stability guarantee.

We cannot present simulations for traditional min–max MPC because the control law is infeasible, and therefore undefined, for the example we consider.

Feedback Min–Max MPC: In feedback min–max MPC, a family of control sequences, $[u^\ell_{t|t}, u^\ell_{t+1|t}, u^\ell_{t+2|t}]$, is optimized at time t, each one corresponding to a different disturbance profile.

In view of the above discussion, we need only consider the extreme disturbance realizations. Over the horizon $N = 3$, these are $\{-1,-1,-1\}$, $\{-1,-1,1\}$, $\{-1,1,-1\}$, $\{-1,1,1\}$, $\{1,-1,-1\}$, $\{1,-1,1\}$, $\{1,1,-1\}$ and $\{1,1,1\}$. The corresponding state trajectories are depicted in Fig. 3, and we have the following predictions:

$$
\begin{aligned}
x^1_{t+1|t} &= x_t + u_{t|t} - 1 & x^1_{t+2|t} &= x_t + u_{t|t} + u^1_{t+1|t} - 2 \\
x^2_{t+1|t} &= x_t + u_{t|t} + 1 & x^2_{t+2|t} &= x_t + u_{t|t} + u^1_{t+1|t} \\
& & x^3_{t+2|t} &= x_t + u_{t|t} + u^2_{t+1|t} \\
& & x^4_{t+2|t} &= x_t + u_{t|t} + u^2_{t+1|t} + 2
\end{aligned}
$$

$$
\begin{aligned}
x^1_{t+3|t} &= x_t + u_{t|t} + u^1_{t+1|t} + u^1_{t+2|t} - 3 \\
x^2_{t+3|t} &= x_t + u_{t|t} + u^1_{t+1|t} + u^1_{t+2|t} - 1 \\
x^3_{t+3|t} &= x_t + u_{t|t} + u^1_{t+1|t} + u^2_{t+2|t} - 1 \\
x^4_{t+3|t} &= x_t + u_{t|t} + u^1_{t+1|t} + u^2_{t+2|t} + 1 \\
x^5_{t+3|t} &= x_t + u_{t|t} + u^2_{t+1|t} + u^3_{t+2|t} - 1 \\
x^6_{t+3|t} &= x_t + u_{t|t} + u^2_{t+1|t} + u^3_{t+2|t} + 1 \\
x^7_{t+3|t} &= x_t + u_{t|t} + u^2_{t+1|t} + u^4_{t+2|t} + 1 \\
x^8_{t+3|t} &= x_t + u_{t|t} + u^2_{t+1|t} + u^4_{t+2|t} + 3.
\end{aligned}
\qquad (18)
$$

Consider now the control sequences, $u^\ell_{j|t}$, defined by

$$
\begin{aligned}
u_{t|t} &= -x_t & u^1_{t+1|t} &= 1 & u^1_{t+2|t} &= 1 \\
& & u^2_{t+1|t} &= -1 & u^2_{t+2|t} &= -1 \\
& & & & u^3_{t+2|t} &= 1 \\
& & & & u^4_{t+2|t} &= -1.
\end{aligned}
\qquad (19)
$$

Fig. 3.   Example: possible state trajectories.

These lead to the following predictions:

$$
\begin{aligned}
x^1_{t+1|t} &= -1 & x^1_{t+2|t} &= -1 & x^1_{t+3|t} &= -1 \\
x^2_{t+1|t} &= 1 & x^2_{t+2|t} &= 1 & x^2_{t+3|t} &= 1 \\
& & x^3_{t+2|t} &= -1 & x^3_{t+3|t} &= -1 \\
& & x^4_{t+2|t} &= 1 & x^4_{t+3|t} &= 1 \\
& & & & x^5_{t+3|t} &= -1 \\
& & & & x^6_{t+3|t} &= 1 \\
& & & & x^7_{t+3|t} &= -1 \\
& & & & x^8_{t+3|t} &= 1.
\end{aligned}
\qquad (20)
$$

For all x_t, we find, therefore, that the above set of controls satisfies the stability and state constraints. The control optimization is therefore feasible for all initial states, and it follows that the optimal feedback min–max MPC law is always defined. Moreover, this control law leads to no constraint violations and is guaranteed stabilizing, as long as $w_t \in [w_-, w_+]$.

To conclude, we present simulation results obtained by use of feedback min–max MPC law. In Fig. 4, we show the results obtained with $x_0 = 2$ and $w_t = -1/t$, $t \geq 1$; in Fig. 5, the results obtained with $x_0 = 2$ and $w_t = -\cos(t/5)$, $t \geq 0$, are presented. As expected, we find in both cases, that in spite of the disturbance, the control law drives the state to the robust invariant set $\mathbf{X}_0$ allowing no constraint violations at any time.

Fig. 4.   Example: min–max MPC, $w_t = -1/t$.

Fig. 5.   Example: min–max MPC, $w_t = -\cos(t/5)$.

III. VARIABLE HORIZON FORMULATION

The use of variable horizons in MPC has recently come under increasing scrutiny, due to the advantages it offers over fixed horizons. First, the use of a variable horizon removes the need to design the fixed horizon, which can be a tricky task. Variable horizon control schemes, moreover, can lead to identical nominal performance in open-loop and receding-horizon implementations; provided there is no model error, this property of the control may be used to avoid any further computation, without loss of performance, once an acceptable initial trajectory has been computed. In practice, there is always some deviation, but the control assuming zero model error is usually a good first guess for further optimization.

In this section, we formulate the variable horizon variant of feedback min–max MPC and discuss its properties.

The variable horizon scheme we consider is based on the following optimization:

$$\min_{\{u_{j|t}^{\ell}\}, N} \max_{\ell \in \mathcal{L}} N$$

$$\text{s.t.} \quad \begin{aligned} x_{j|t}^{\ell} &\in \mathbf{X}, & j &> t, \quad \forall \ell \in \mathcal{L}, \\ u_{j|t}^{\ell} &\in \mathbf{U}, & j &\geq t, \quad \forall \ell \in \mathcal{L}, \\ x_{t+N|t}^{\ell} &\in \mathbf{X}_0, & & \forall \ell \in \mathcal{L}, \\ x_{j|t}^{\ell_1} = x_{j|t}^{\ell_2} &\Rightarrow u_{j|t}^{\ell_1} = u_{j|t}^{\ell_2}, & j &\geq t, \quad \forall \ell_1, \ell_2 \in \mathcal{L} \end{aligned} \quad (21)$$

where the control horizon N becomes a degree of freedom. The outer controller is obtained by receding-horizon implementation of the solution to this optimization. As before, the inner controller, $u = -Kx$, is used when the state enters $\mathbf{X}_0$.

A. The Control Algorithm

The variable horizon feedback min–max MPC law we propose is summarized as follows.

Algorithm 2—Variable Horizon Feedback Min–Max MPC:

Data: x_t

Algorithm: If $x_t \in \mathbf{X}_0$, set $u_t = -Kx_t$. Otherwise, find the solution of (21) and set u_t to the first control in the optimal sequence calculated.

B. Properties

Because the cost in the variable horizon control law is the time taken to reach the robust control invariant set $\mathbf{X}_0$ the controller drives the state to this set in minimal time. Stability is also easily established.

Theorem 3: The variable horizon feedback min–max MPC law given by Algorithm 2 drives the state x_t to the robust control invariant set $\mathbf{X}_0$ in finite time and keeps it in this set for all subsequent times.

Proof: Let $\{w_{j|t}^{\ell}\}$, $\ell \in \mathcal{L}$ represent the possible disturbance realizations, and let $\{\bar{u}_{j|t}^{\ell}\}$ denote the corresponding optimal control sequences at time t. Further, given that the state at time t is x_t, let $\{x_{j|t}^{\ell}\}$ and ϕ_t^{ℓ} denote the state trajectories and costs associated with each of the optimal control/disturbance realizations. The optimal cost, at time t, is $\bar{\phi}_t = \max_{\ell \in \mathcal{L}} \phi_t^{\ell}$. At time t, the first of the optimal controls is applied and the disturbance takes a certain value w_t. Let $\mathcal{L}_1$ denote the set of indexes such that $w_{t|t}^{\ell} = w_t$ for all $\ell \in \mathcal{L}_1$ and $w_{t|t}^{\ell} \neq w_t$ for all $\ell \notin \mathcal{L}_1$. At time $t+1$, the state has moved along a trajectory that coincides with the predicted state trajectories indexed by $\ell \in \mathcal{L}_1$. The control sequences $[\bar{u}_{t+1|t}^{\ell}, \bar{u}_{t+2|t}^{\ell}, \cdots, \bar{u}_{t+N-1|t}^{\ell}, -Kx_{t+N|t}^{\ell}]$, $\ell \in \mathcal{L}_1$ then satisfy the constraints and yield costs $\phi_t^{\ell} - 1$. The optimal cost at time $t+1$ is no larger than largest of these costs; denoting by $\bar{\phi}_{t+1}$ the optimal cost at time $t+1$, we therefore get, $\bar{\phi}_{t+1} \leq \max_{\ell \in \mathcal{L}_1} \phi_t^{\ell} - 1$, and therefore

$$\bar{\phi}_{t+1} \leq \bar{\phi}_t - 1 \quad (22)$$

since $\max_{\ell \in \mathcal{L}_1} \bar{\phi}_t^{\ell} \leq \max_{\ell \in \mathcal{L}} \bar{\phi}_t^{\ell} = \bar{\phi}_t$. The cost therefore decreases to zero in finite time. As $\bar{\phi}_t > 0$ for all $x_t \notin \mathbf{X}_0$, it follows that the

state enters $\mathbf{X}_0$ in finite time. Finally, once the state enters $\mathbf{X}_0$, it is subsequently kept inside this set by the inner controller $u = -Kx$. $\square$

Like Theorem 1, for the fixed horizon case, this theorem establishes stability only for the case where the disturbance follows one of the realizations considered in the optimization. If the disturbance veers away from the realizations considered, the stability proof no longer applies. However, as in the fixed horizon case, linearity of the process and convexity of the constraints allow us to obtain a stability result that states that only the extreme realizations of the disturbance need to be considered in the min–max optimization in order to obtain a control formulation that is stabilizing for all disturbance profiles that lie in the convex hull of the extreme realizations, i.e., all disturbance sequences that take values in $\mathbf{W}$ (see Remark 3).

As before, we assume $\mathbf{W}$ is a polytope in $\mathbf{R}^n$, and we let $\mathcal{L}_v$ denote the set of indexes ℓ, such that $\{w_{j|t}^\ell\}$ takes values only on the vertices of $\mathbf{W}$. We then consider the min–max control optimization

$$\min_{\{u_{j|t}^\ell\}, N} \max_{\ell \in \mathcal{L}_v} N$$

$$\text{s.t.} \quad \begin{aligned} x_{j|t}^\ell &\in \mathbf{X}, & j &> t, & \forall \ell &\in \mathcal{L}_v, \\ u_{j|t}^\ell &\in \mathbf{U}, & j &\geq t, & \forall \ell &\in \mathcal{L}_v, \\ x_{t+N|t}^\ell &\in \mathbf{X}_0, & & & \forall \ell &\in \mathcal{L}_v, \\ x_{j|t}^{\ell_1} &= x_{j|t}^{\ell_2} \Rightarrow u_{j|t}^{\ell_1} = u_{j|t}^{\ell_2}, & j &\geq t, & \forall \ell_1, \ell_2 &\in \mathcal{L}_v \end{aligned} \quad (23)$$

and we obtain the outer control by receding-horizon implementation of the solution to this optimization. Again, this optimization has finite dimension, because $\mathcal{L}_v$ contains only a finite number of indexes. We show that the resulting min–max control law satisfies the state and input constraints and is stabilizing for all disturbance realizations $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}$.

Theorem 4: Let $\mathbf{W}$ be a polytope in $\mathbf{R}^n$, with $\{w_{j|t}^\ell\}$, $\ell \in \mathcal{L}_v$ denoting the extreme disturbance realizations, which take values at the vertices of $\mathbf{W}$. Then, the feedback min–max law given by Algorithm 2, with the optimization of (21) replaced by (23), drives the state x_t to the robust control invariant set $\mathbf{X}_0$ in finite time and keeps it in this set for all subsequent times.

Proof: The first part of the proof proceeds exactly as the proof for Theorem 2; we find that, at time $t + 1$, state x_{t+1}, the control strategy of (10) leads to satisfaction of the state, input and stability constraints. Also, because under that control strategy the state is forced to $\mathbf{X}_0$ in one less sample than at time t, we find that the control sequence of (10) yields a cost $\bar{\phi}_t - 1$, where $\bar{\phi}_t$ denotes the optimal cost at time t. The optimal cost at time $t + 1$ is therefore no larger than $\bar{\phi}_t - 1$, and we find consequently that the cost decreases to zero in finite time. As the cost is greater than zero for all states outside the invariant set $\mathbf{X}_0$ this leads to the conclusion that the state enters $\mathbf{X}_0$ in finite time. Finally, once the state enters $\mathbf{X}_0$, it is subsequently kept inside this set by the inner controller $u = -Kx$. $\square$

C. Computational Issues

Because of the simplicity of the cost, which involves only the horizon N, the formulation of the control optimization is simpler here than in the fixed horizon case.

The control optimization for the variable horizon formulation of the control law is a mixed-integer nonlinear program, of dimension $m(1 + p + p^2 + \cdots + p^{N-1}) + 1$, the extra degree of freedom being due to the addition of the variable horizon N as a degree of freedom. However, as the cost involves only the integer-valued variable N, the efficient solution of the optimization results from a simple integer search for the lowest N for which there exists a control sequence that satisfies the constraints. Once the optimal trajectory has been calculated at time t, the optimal trajectory at all subsequent times is known, provided there is no model error. This is because

implementation of the appropriate linear combination of the control sequences calculated at time t satisfies the constraints and yields a cost reduction of one at each sample until the state enters $\mathbf{X}_0$ (see proofs of Theorems 2 and 4). This performance must be optimal since it cannot be improved upon if, indeed, the trajectory calculated at time t was optimal. When the process model is exact, and the disturbance remains in $\mathbf{W}$, therefore, the control optimization needs to be performed only once.

IV. CONCLUSION

In this paper, we have outlined the details of min–max MPC formulations that introduce, in the control optimization, the notion that feedback is present in the control implementation. This often leads to improved performance compared to standard MPC schemes, which do not take into account the effects of possible future disturbances. The method also avoids the likely feasibility problems that result from the use of min–max formulations that optimize a single control profile over all possible future disturbance realizations. These points were illustrated with a simple example.

The price that must be paid for these benefits is that the computational demands of the feedback min–max algorithms can be very high, both in the fixed and variable horizon cases. We showed that because the process is linear, all possible disturbance realizations do not need to be considered in the optimization; the control needs to be optimized only over the extreme disturbance realizations. However, the number of extreme realizations grows combinatorially with the horizon N. Although the number of control profiles that are computed at each sample is smaller than the number of disturbance realizations, due to the causality constraint, this can make the control optimization prohibitively expensive if large horizons are used. The method can, however, be very effective with small horizons. Furthermore, when the process model is exact and the disturbance remains in $\mathbf{W}$, the control optimization needs to be performed only once, and reoptimization is not necessary at each sample. Hence, if the model error is small, the current control differs little from that obtained from the previous optimization which, therefore, provides an excellent initial guess.

REFERENCES

[1] 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.

[2] J. B. Rawlings and K. R. Muske, "Stability of constrained receding horizon control," *IEEE Trans. Automat. Contr.*, vol. 38, no. 10, pp. 1512–1516, Oct. 1993.

[3] J. B. Rawlings, E. S. Meadows, and K. R. Muske, "Nonlinear model predictive control: A tutorial and survey," in *Proc. ADCHEM'94*, Kyoto, Japan.

[4] H. Genceli and M. Nikolaou, "Robust stability analysis of constrained l_1-norm model predictive control," *AIChE J.*, vol. 39, no. 12, pp. 1954–1965, 1993.

[5] Z. Q. Zheng and M. Morari, "Robust stability of constrained model predictive control," in *Proc. 1993 American Control Conf.*, pp. 379–383.