ADCHEM 2006

International Symposium on Advanced Control of Chemical Processes
Gramado, Brazil – April 2-5, 2006

IFAC

NONCONVEX OPTIMIZATION AND ROBUSTNESS IN REALTIME MODEL PREDICTIVE CONTROL

Darryl DeHaan, [1] **Martin Guay** [2]

Department of Chemical Engineering, Queen's University, Kingston, ON, Canada K7L 3N6

Abstract: Recent works in the nonlinear MPC literature have presented "realtime" optimization approaches based upon incremental updating of input parameters using local descent directions of the cost functional. The main downside to these methods is their strong dependence upon the values used to initialize the input parameters. In this note we study the robustness issues associated with non-local search methods in continuous-time MPC, and demonstrate a framework for robustly incorporating these approaches in a realtime setting.

Keywords: nonlinear model predictive control, robustness, nonlinear systems, realtime optimization

1. INTRODUCTION

The flexibility of model predictive control (MPC) for dealing with constraints has lead to its rapid emergence as the advanced control method of choice in the process industries. However, computational complexity remains the main limitation preventing the use of MPC in many applications.

Current application of nonlinear MPC is limited to so-called "perfect model" implementations, which rely on nominal robustness guarantees such as discussed in (Magni and Sepulchre, 1997; Grimm *et al.*, 2004; Grimm *et al.*, 2003). In particular, in (Grimm *et al.*, 2004) it is shown that continuity of *either* the value function *or* the MPC feedback policy are sufficient for nominal robustness to disturbance inputs; while the feedback policy may generally be discontinuous for nonlinear problems, the value function can be made continuous using appropriate inner approximations of the constraint limits.

However, as discussed in (Coron and Rosier, 1994), discontinuous feedback policies for continuous-time systems are potentially non-robust to measurement error in the feedback loop. While (Messina *et al.*, 2005) show that *discrete-time* MPC exhibits nominal robustness to measurement noise, (Tuna *et al.*, 2005) demonstrate that this nominal robustness may approach zero for systems with fast sampling.

In this work, we demonstrate that realtime MPC methods based on gradient-driven local optimization are automatically nominally robust to measurement errors. Furthermore, we proceed to demonstrate a means by which non-local optimization methods can be incorporated into a realtime framework without violating nominal robustness. This work is organized as follows. Section 2 discusses the lack of robustness in the standard definition of MPC, while Section 3 reviews the basic ideas of realtime MPC and presents robustness results. Section 4 discusses the incorporation of nonconvex optimization into a realtime framework, and a simulation example is included in Section 5.

---

[1] Work supported by the National Sciences and Engineering Research Council of Canada.
[2] Corresponding author. email: guaym@chee.queensu.ca.

1.1 Preliminaries

Throughout this work, $\|s\|_\infty$ denotes a vector ∞−norm, whereas the space $\mathcal{L}^\infty_\mathcal{I}$ of bounded functions on domain $\mathcal{I}$ have norm $\|s(\cdot)\|_{\mathcal{L}^\infty}$. A function $\gamma : [0, \infty) \to \mathbb{R}_{\geq 0}$ is of class $\mathcal{K}$ if it is monotone increasing from $\gamma(0) = 0$. The notations $\overline{\mathbb{S}}$, $\mathring{\mathbb{S}}$, $\partial\mathbb{S}$, co$\{\mathbb{S}\}$, and $\mu(\mathbb{S})$ respectively denote the closure, interior, boundary (i.e. $\overline{\mathbb{S}} \setminus \mathring{\mathbb{S}}$), convex hull, and Lebesgue measure of a set $\mathbb{S}$. We define the distance $d_\mathbb{S}(s) \triangleq \inf_{s' \in \mathbb{S}} \|s - s'\|$ and the ball $\mathcal{B}(\mathbb{S}, \varepsilon) \triangleq \{s \mid d_\mathbb{S}(s) \leq \varepsilon\}$. Finally, a function $f \in C^{p+}$ if $f \in C^p$, with all $\nabla^p f$ locally Lipschitz.

The system of interest is any nonlinear dynamic

$$\dot{x} = f(x, u), \qquad x(0) = x_0 \tag{1}$$

subject to pointwise constraints of the form $(x, u) \in \mathbb{X} \times \mathbb{U} \subseteq \mathbb{R}^n \times \mathbb{R}^m$, such that $\mu(\mathbb{X} \times \mathbb{U}) > 0$. The control objective is regulation of x to a (not necessarily connected) target set $\Sigma_x \subset \mathbb{X}$, which is weakly invariant under (1) for $(x, u) \in \Sigma \triangleq \Sigma_x \times \Sigma_u(x)$. Performance is measured by the function

$$J(x, \mathbf{u}_{[0,T]}) = \int_0^T L(x^p, u)\, d\tau + W(x^p(T)) \tag{2a}$$

$$s.t.\quad \dot{x}^p = f(x^p(\tau), u(\tau)), \quad x^p(0) = x \tag{2b}$$

$$(x^p(\tau), u(\tau)) \in \mathbb{X} \times \mathbb{U}, \quad \forall \tau \in [0, T] \tag{2c}$$

$$x^p(T) \in \mathcal{X}_f \tag{2d}$$

Unless stated otherwise, the functions $f(\cdot, \cdot)$, $L(\cdot, \cdot)$, and $W(\cdot)$ are assumed C^{0+}, and there exists $\gamma_1, \gamma_2 \in \mathcal{K}$ such that $L(x, u) \geq \gamma_1(d_\Sigma(x, u))$ and $W(x) \geq \gamma_2(d_{\Sigma_x}(x))$. It is assumed that $\mathcal{X}_f$ and W satisfy sufficient conditions for stability as detailed in (Mayne *et al.*, 2000), and x_0 is feasibly open-loop stabilizable to the origin.

2. CONTINUOUS-TIME MPC AND MEASUREMENT ERRORS

A generic continuous-time MPC feedback $u = k_{mpc}(x)$ which minimizes (2) is given by

$$k_{mpc}(x) = \lim_{\tau \downarrow 0} \mathbf{u}^*_{[0,T]}(\tau) \tag{3a}$$

$$\mathbf{u}^*_{[0,T]} = \arg \min_{\mathbf{u}_{[0,T]}} J(x, \mathbf{u}_{[0,T]}) \tag{3b}$$

$$J^*(x) \triangleq J(x, \mathbf{u}^*_{[0,T]}) \; . \tag{3c}$$

The minimization (3b) is over piecewise continuous functions $\mathbf{u}_{[0,T]} \in \mathcal{L}^\infty_{[0,T]}$, and solutions are not necessarily unique, so $k_{mpc} : \mathbb{X} \to \mathbb{U}$ is a possibly discontinuous and set-valued mapping.

The closed-loop dynamics therefore properly take the form of the differential inclusion

$$\dot{x} \in f(x, k_{mpc}(x)) \; . \tag{4}$$

A uniformly continuous function $x(t)$ is a *classical solution* on the interval $t \in [0, T)$ if it satisfies (4) for almost all $t \in (0, T)$. Traditional MPC stability proofs, such as (Mayne *et al.*, 2000; Chen and Allgower, 1998), correspond to showing $J^*(x)$ is nonincreasing over all classical solutions to (4).

It follows from (Grimm *et al.*, 2003) that if the state constraint in (2) is replaced with one of the form $x(\tau) \in \mathbb{X}'(\tau)$, where $\mathcal{X}_f \subset \mathbb{X}'(t_2) \subset \mathbb{X}'(t_1) \subset \mathbb{X}$ for all $0 \leq t_1 < t_2 \leq T$, then there exists a $\delta > 0$ such that for $\|d\|_{\mathcal{L}^\infty} \leq \varepsilon$, global asymptotic stability (GAS) of classical solutions to (4) implies the same for

$$\dot{x} \in f(x, k_{mpc}(x)) + d(t) \, . \tag{5}$$

In contrast, arbitrarily small measurement error

$$\dot{x} \in f(x, k_{mpc}(x + e(t))) \; . \tag{6}$$

can cause k_{mpc} to "dither" around a discontinuity, generating trajectories which may differ greatly from any classical solution of (4). This illustrates the notion of *Filippov* solutions to (4), defined as any trajectory $x(t)$ satisfying $\dot{x} \in \mathcal{F}(x)$, $\mathcal{F}(x) \triangleq \text{co}\{\overline{f(x, k_{mpc}(x))}\}$, for almost all $t \in (0, T)$. This type of dithering can induce directions of motion not even in the span of $\nabla_u f(x, \cdot)$, and hence not covered by the standard stability proof for k_{mpc}.

To demonstrate the potential impact of measurement noise, in similar spirit to (Tuna *et al.*, 2005) we offer the following extension of (Coron and Rosier, 1994, Prop. 1.4). First, it is necessary to (loosely) define the *Clarke normal cone* $N^C_\mathbb{S}(s)$ at $s \in \partial\mathbb{S}$ as the convex hull of every vector normal to $\mathbb{S}$ at arbitrary s', with $s' \to s$ (Clarke *et al.*, 1998).

Lemma 1. Assume all classical solutions to (4) are GAS. Define $\mathcal{H} = \{x \in \mathbb{X} \,|\, \overline{k_{mpc}(x)} \setminus k_{mpc}(x) \neq \emptyset\}$, and for each $z_0 \in \mathcal{H}$, let $\mathcal{Z}(z_0)$ denote the set of Filippov solutions $z : \mathcal{I} \to \mathbb{R}^n$ to (4) with $z(0) = z_0$ and $\mathcal{I} \subseteq [0, \infty)$ maximal. Then any initial state $x(0) \in \mathcal{H}_u$

$$\mathcal{H}_u \triangleq \{z_0 \in \mathcal{H} \,|\, \sup_{z(\cdot) \in \mathcal{Z}(z_0)} \lim_{t \uparrow \partial\mathcal{I}} d_\Sigma(z(t)) > 0\}$$

can be prevented from reaching Σ_x by arbitrarily small measurement error. Furthermore, if the set

$$\mathcal{H}^o_u \triangleq \{z_0 \in \mathcal{H}_u \,|\, \forall \xi \in N^C_{\mathcal{H}_u}, \exists \nu \in \mathcal{F}(z_0)\; s.t.\; \langle \xi, \nu \rangle > 0\}$$

is nonempty, then $\exists \delta(\cdot) \in \mathcal{K}$ and a neighbourhood $\mathfrak{H}$ of $\mathcal{H}^o_u$ with measure $\mu(\mathfrak{H}) \geq \delta(\varepsilon) > 0$ such that any initial state $x(0) \in \mathfrak{H}$ can be prevented from reaching Σ by a measurement error $\|e\|_{\mathcal{L}^\infty} \leq \varepsilon$.

The first claim is essentially a direct application of (Coron and Rosier, 1994, Prop. 1.4). The second claim is an extension, which essentially states that if Filippov solutions span every outward direction, then measurement error can induce flows back *towards* $\mathcal{H}^o_u$ from any point sufficiently close (i.e.

with x and $x+e$ on "opposite sides" of $\mathcal{H}_u^o$). This is relatively straightforward given $f \in C^{0+}$.

Within the context of chemical processes, there is potential for this type of effect any time a controller must make a choice between two distinct paths. One major instance in which this occurs is when x is near the threshold at which a profitable trajectory becomes infeasible with respect to state constraints, and must be abandoned for a less profitable one.[3] Another situation of interest is when pieces of equipment must be arbitrarily selected from a group for some type of preferential treatment (e.g. starting up parallel pumps in sequence to avoid electrical trips, applying activation heat to parallel reactors sequentially due to steam limitations, etc). This is of particular concern when responding to unplanned events requiring decisive response to mitigate losses.

3. DESCENT-BASED REALTIME METHODS

It is apparent from Section 2 that robustness issues due to measurement noise stem from the assumed globality of the minimization in (3). In practice, however, only local solutions can be guaranteed online; in fact, precisely locating even a local minimum within one sampling period can be difficult. For this reason, a variety of "real-time" approaches such as (Ohtsuka, 2004; Cannon and Kouvaritakis, 2000; DeHaan and Guay, 2005) allow the optimization parameters to evolve incrementally within the same timescale as the dynamics. In this section, we study whether these approaches suffer the same lack of robustness to measurement error discussed in Section 2. Our presentation will follow the method of (DeHaan and Guay, 2005), reviewed briefly,[4] but an effort is made to generalize whenever possible.

3.1 Description of Realtime Method

The approach in (DeHaan and Guay, 2005) allows for piecewise parameterization of the control input using a parameter vector $\omega^T = [\pi; \theta]$, consisting of an ordered time support $\pi \in \mathbb{R}^{N+1}_{\geq 0}$ and vectors $\theta_i \in \mathbb{R}^p$ for each of the N intervals. Using a preselected basis ϕ (e.g. polynomial, exponential), the input trajectory is parameterized piecewise as

$$u^p(\tau,\omega) = \begin{cases} \phi(\tau,\theta_1) & \tau \in [0,\, \pi_1] \\ \phi(\tau - \pi_{i-1},\, \theta_i) & \begin{array}{c}\tau \in (\pi_{i-1},\, \pi_i], \\ i \in \{2 \dots N\}\end{array} \end{cases} \quad (7)$$

which is substituted for $u(\tau)$ in (2), to define $J(x,\omega)$ in an obvious manner. The closed-loop system evolves continuously as

$$\begin{aligned} \dot{x} &= f(x, k_{rt}(\omega)), \quad k_{rt}(\omega) = u^p(\pi_0, \theta_1) && (8a) \\ \dot{\omega} &= g(x,\omega) && (8b) \\ & \text{while } h(\omega) \triangleq \pi_1 - \pi_0 > 0 && (8c) \end{aligned}$$

where $\omega(0) = \omega_0$ is assumed to be chosen as a feasible (but sub-optimal) parameterization with respect to (2b,c). Upon equality $h(\omega) = 0$ (i.e. the first switching point of (7) is reached), the reset

$$\omega^+ = \begin{cases} \pi_i^+ = \begin{cases} 0 & i = 0 \\ \pi_{i+1} & 1 \leq i < N \\ \pi_N + \delta(x^p(\pi_N)) & i = N \end{cases} \\ \theta_i^+ = \begin{cases} \theta_{i+1} & 1 \leq i < N \\ \kappa(x^p(\pi_N)) & i = N \end{cases} \end{cases} \quad (9a)$$

$$\text{when } h(\omega) \triangleq \pi_1 - \pi_0 = 0 \quad (9b)$$

occurs, where $\delta : \mathcal{X}_f \to \mathbb{R}_{>0}$ and $\kappa : \mathcal{X}_f \to \mathbb{R}^p$ are C^{0+} functions representing a local control law which feasibly initializes the parameters for the new interval being augmented to the horizon. This makes the closed-loop behaviour that of a hybrid system[5]. Stability is proven by the invariance principle, using the strict decrease of $J(x,\omega)$ under (8), the non-increase under (9)[6], and the fact that $h(\omega) \equiv 0$ is not invariant.

The precise definition of $g(x,\omega)$ in (8) varies somewhat between methods, in particular depending on the manner by which the constraints (2c,d) are enforced. We thus distinguish between two classes of approach. In all cases, gradients of J are ensured to exist by strengthening the assumptions on $L(\cdot,\cdot)$, $W(\cdot)$, $f(\cdot,\cdot)$ to C^{1+}.

3.1.1. Interior-point Approaches In the approach of (DeHaan and Guay, 2005), the constraints (2c,d) are replaced by an augmented cost $J^a(x,\omega)$ which incorporates C^{1+} barrier functions into $L(x,u)$ and $W(x)$. Then $g(x,\omega)$ is of the form

$$g(x,\omega) = \mathcal{G}(\omega, \upsilon), \quad \upsilon \triangleq -k\nabla_\omega J^a(x,\omega) \quad (10)$$

where υ represents the nominal descent-update, while $\mathcal{G}$ is a locally Lipschitz operator, possibly including such operations as projecting υ to keep ω in a desired convex set, limiting the growth rate of ω, etc. In particular, $\mathcal{G}$ ensures both $\langle \nabla_\omega J^a, \mathcal{G}(\omega,\upsilon)\rangle \leq 0$ and $\lim_{h(\omega)\downarrow 0}(\dot{\pi}_1 - \dot{\pi}_0) < 0$ (so event (9b) is well defined under small perturbations). Stability follows using J^a as a Lyapunov function, since $\frac{dJ^a}{dt} \leq -L(x, k_{rt}(\omega)) + \langle \nabla_\omega J^a, \mathcal{G}(\omega,\upsilon)\rangle < 0$. The open interior of all constraints is rendered invariant, thus preserving feasibility given feasible initial conditions.

[3] this is *still possible* even if (2b) is nominally robust with respect to $d(t)$, and $\mathbb{X}(\tau)$ is nested as discussed previously.
[4] for omitted details, see (DeHaan and Guay, 2005)

[5] While (Ohtsuka, 2004; Cannon and Kouvaritakis, 2000) do not include hybrid behaviour, they are still encompassed in (7)-(9) by simply defining $\dot{\pi} \equiv 1$, $\pi_0(0) \neq \pi_1(0)$.
[6] guaranteed by conditions on the design of $\kappa(x)$ and $\delta(x)$

3.1.2. Active set-based Approaches In the approach of (Cannon and Kouvaritakis, 2000), the update law is of the form (10), but with $\nabla_\omega J$ instead of $\nabla_\omega J^a$. Constraint feasibility is ensured by incorporation a parameter projection into the definition of $\mathcal{G}$, which removes components of v directed out of the (x-dependent) feasible set for ω. Since the feasible parameter set is impractical to calculate online, this approach corresponds to testing along the length of the prediction $x^p(\tau)$ for constraint violation. Effectively, this is comparable to incorporating a lagrange-multiplier trajectory $\lambda(\tau)$ into the minimization. For robustness, this approach must include a τ-varying constraint $\mathbb{X}'(\tau)$ of the form discussed in Section 2.

3.2 Robustness of Descent-based Methods

The question of robustness for the class of descent-based approaches discussed in Section 3.1 is answered by the following Lemma. Measurement errors in (8)-(10) are interpreted as (2) taking the form $J(x + e(t), \omega)$ when evaluating $\nabla_\omega J$ and $x^p(\pi_N)$ in the feedbacks.

Lemma 2. Assume $f(\cdot,\cdot), L(\cdot,\cdot), W(\cdot) \in C^{1+}$. For any initial condition (x_0, ω_0) such that predictions (2b-d) are strictly feasible, global asymptotic stability of the target Σ_x under closed-loop dynamics (8)-(9) is nominally robust to additive disturbances and measurement errors.

Sketch of Proof: By standard results, the prediction $x^p(\tau)$ of (2b) is Lipschitz with respect to small changes in its initial condition x. Then for interior-point approaches, $\nabla_\omega J^a(x, \omega)$ is C^{0+} in x, as are the dynamics of (8), (9) and the jump condition $h(\omega)$. If $z(t)$ is any closed-loop solution from $z(0) = x(0)$, perturbed by $\|d, e\|_{\mathcal{L}^\infty} \le \varepsilon$, then it follows that $x(t)$ and $z(t)$ are "close" in the sense $\lim_{\varepsilon \downarrow 0}\{\max_{s\in\mathbb{R}^+} \min_{s'\in\mathbb{R}^+} \|x(s) - z(s')\|\} = 0$. Feasibility and approximate-GAS of $z(\cdot)$ for sufficiently small $\varepsilon > 0$ follow by the *strict* feasibility and asymptotic stability of $x(\cdot)$. A similar argument works for active-set approaches. ■

The robustness margins discussed above may be quite small, since the proof makes no use of the update (8b) to compensate by adjusting the input parameterization (since it cannot be guaranteed that this is possible, when $\nabla_\omega J \approx 0$). However, it is important to note that any method based upon achieving minimizing solutions (as opposed to incremental improvement) must necessarily be implemented in a sampled-data framework, with significantly slower sampling rates than used by realtime methods. Such methods then require an equally restrictive assumption of open-loop robustness between sampling instances.

4. ROBUSTLY INCORPORATING NONCONVEX METHODS

Since descent-based realtime methods evolve continuously within connected neighbourhoods of their initial conditions, they do not generate differential inclusions (assuming C^{1+} process dynamics). As such, a local descent-based method automatically exhibits nominal robustness to measurement errors, unlike global approaches. The obvious downside is that performance depends strongly upon the quality of the initial condition.

One simple approach to reducing the impact of ω_0 is to instead allow for multiple initial conditions to be specified; i.e. to allow for multiple independent sets of parameters to be (locally) minimized in parallel. It will first be shown how this can be done in a manner robust to measurement noise, after which we will provide discussion on more sophisticated extensions of this approach.

4.1 Modified Realtime Dynamics

Let $q \ge 1$ denote the number of independent input parameter sets calculated online, denoted ω^i, $i \in \{1, \dots q\}$. For convenience it is assumed all parameterizations use the same basis ϕ, although each ω^i contains an independent time support π^i, potentially of length $N^i \ne N^j$ intervals. The flows (8) then take the form

$$\dot{x} = f(x, k_{rt}(\omega^{i^*})), \;\; k_{rt}(\omega^{i^*}) = u^p(\pi_0^{i^*}, \theta_1^{i^*}) \quad (11a)$$

$$\dot{\omega}^i = \begin{cases} g^i(x, \omega^i) & i = i^* \\ \mathrm{g}^i(x, \omega^i) & i \ne i^* \end{cases} \quad \forall i \in \{1 \dots q\} \quad (11b)$$

$$\text{while} \;\; \pi_1^{i^*} - \pi_0^{i^*} > 0, \quad \text{and} \quad (11c)$$

$$\text{while} \;\; J(x, \omega^{i^*}) < (1+\epsilon) J(x, \omega^i), \; \forall i \ne i^* \quad (11d)$$

where $\epsilon > 0$ is a design constant, and $i^* \in \mathbb{N}$ is an additional state of the system identifying which one of the parameter sets is "active". The vectorfields $g^i(x, \omega^i)$ are of the form discussed in Section 3, while g^i reflects that a different update law may be desired if ω^i is not active.

Equality of (11c) triggers the reset (9) (for ω^{i^*} only). Meanwhile, equality of (11d) triggers

$$(i^*)^+ = \text{choose}\left\{ \arg \min_{i\in\{1\dots q\}} J(x, \omega^i) \right\} \quad (12)$$

in which "choose" represents any arbitrary convention for selecting between multiple minimizers. If an interior-point approach is used for the $g^i(x, \omega^i)$, then (11d) and (12) are interpreted using $J^a(x, \omega^i)$; regardless, we adopt the convention that $J(x, \omega^i) = +\infty$ if either (2c,d) do not hold.

Using $J(x, \omega^{i^*})$ as a Lyapunov function, the previous stability arguments still hold in light of the strict decrease of $J(x, \omega^{i^*})$ under reset (12). The robustness of this stability to small measurement errors $e(t)$ hinges on the fact that $\epsilon > 0$ in (11d); i.e. for $\epsilon = 0$, (11a) becomes a differential inclusion with $u = k_{rt}(\Omega)$, $\Omega = \{\omega^i \,|\, J(x,\omega^i) = \min_j J(x,\omega^j)\}$, and is no longer robust to arbitrarily small measurement errors. The nominal robustness of (11) can therefore be summarized:

Claim 3. $\exists\varepsilon(\cdot) \in \mathcal{K}$ s.t. (11) is robust to measurement errors $\|e\|_{\mathcal{L}^\infty} \le \varepsilon(\epsilon)$, with ϵ from (11d)

4.2 Possible Extensions

It is crucial to note that stability depends upon feasible update and resets of the *active* parameterization ω^{i^*} only. Infeasibility of any w^i, $i \neq i^*$, simply results in that parameterization being (temporarily) excluded from consideration as a candidate in (12). Likewise, arbitrary reset mappings of the form $(\omega^i)^+ = \mathfrak{h}^i(x, \omega^1, \ldots, \omega^p)$, $i \neq i^*$ can be executed at any time without impacting stability. To discuss the wide variety of optimization approaches that could be used to design g^i and $\mathfrak{h}^i$ is beyond the scope of this note. Instead, we simply highlight a few possible approaches that could be used,[7] and postpone analysis of their relative merits to future research.

Infeasible-point approaches: The likelihood of identifying the global (or at least a better) minimum significantly increases if the optimization is allowed to temporarily pass through infeasible regions. Thus, the $\mathrm{g}^i(x,\omega^i)$ update laws may be based upon infeasible-point methods without any weakening of guaranteed feasibility of $x(t)$.

Quasi-Global Search Methods: Several methods exist for generating a continuous search trajectory which attempts to visit many or all of the minima on the surface. This includes augmenting the search with extra velocity dynamics to help escape shallow minima, or methods which switch between ascent and descent phases. The $\mathrm{g}^i(x,\omega^i)$ could be based on any one of these.

Deterministic Resetting If $\left\|\nabla_{\omega^i} J(x,\omega^i)\right\|$, $i \neq i^*$, becomes sufficiently small without triggering (11d), then it may be desirable to reset ω^i. One could, for example, generate a *very* crude branch-and-bound (off-line, or in a slower timescale) to identify new values for ω^i showing potential as minimizers, differing sufficiently from other ω^j.

Stochastic Resetting Similar to the above, there are many stochastic approaches which could be used to reset ω^i. This can be done in conjunction with a deterministic master, which shapes the statistical distribution to target regions of interest or rule out regions which are not.

[7] references regrettably omitted for space considerations

5. SIMULATION EXAMPLE

In order to illustrate the approach, we consider a simple exothermic reaction $A \to B$ taking place in a non-isothermal, gas-phase CSTR. The system is comprised of three states; although many equivalent coordinate systems can be used, the equations are most clear for the choice: n (total moles of gas in reactor), n_A (moles of A), and T (reactor temperature). The control objective is specified as regulation to the target $(n, n_A, T)_{ss} = (2.5kmol, 0.25kmol, 500K)$, corresponding to an ideal gas pressure of $1040kPa$, from the initial state $(n, n_A, T)_0 = [2, 0.5, 350]$. Manipulated variables are the outlet molar flow F_{out}, and rate of heat removal $\dot{Q}$. Using several simplistic assumptions, the system equations are

$$\dot{n} = F_{in} - F_{out} \tag{13a}$$

$$\dot{n}_A = F_{in} - \frac{n_A}{n} F_{out} - k(T) n_A \tag{13b}$$

$$\dot{T} = F_{in} \frac{T_{in} - T}{n} - k(T) \frac{\Delta H_r}{\overline{c_p}} \frac{n_A}{n} + \frac{\dot{Q}}{\overline{c_p}\, n} \tag{13c}$$

where $k(T) = k_0 e^{-\frac{E}{RT}}$. System parameters are

$$\Delta H_r = -5000 \frac{kJ}{kmol} \qquad R = 8.314 \frac{m^3\, kPa}{kmol\, K}$$

$$V = 10m^3 \quad E = 8000 \frac{m^3\, kPa}{kmol} \quad \overline{c_p} = 10 \frac{kJ}{kmol\, K}$$

$$T_{in} = 300K \quad k_0 = 6.2s^{-1} \quad F_{in} = 0.25kmol\, s^{-1}$$

An important objective of the controller is to ensure the system trajectories avoid passing through the shaded region in the $P-T$ plane shown in Figure 1, for example to avoid undesireable thermodynamic behaviour of other components in the gas stream which occurs in that region. Using the ideal gas law, this constraint was incorporated as (2c) in the form $g(n,T) \le 0$.

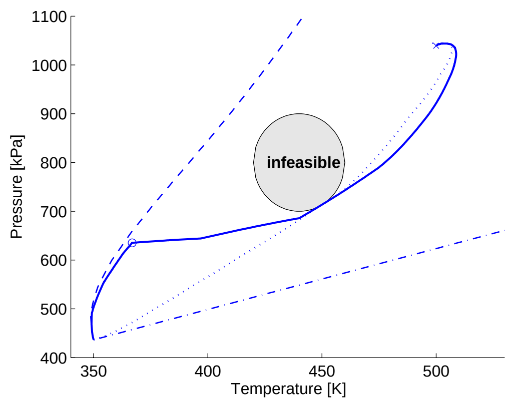

Fig. 1. System trajectories in the $P-T$ plane: closed loop (solid), optimal (dots), initializations ω_0^1 (dash), ω_0^2 (dash-dot). Small circle indicates switch in active ω^i.

The cost function was assumed to be $L(x,u) = x^T \mathrm{diag}(10,\, 10,\, 1\text{e-}3)x + u^T \mathrm{diag}(0.2,\, 0.5)u$, where

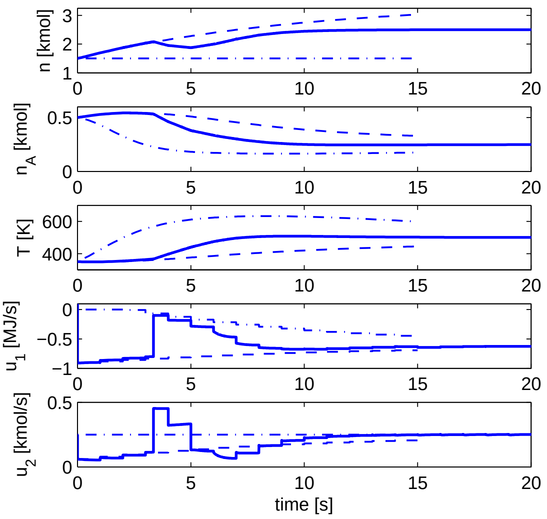

Fig. 2. System trajectories: closed loop (solid), initializations ω_0^1 (dash), ω_0^2 (dash-dot)

$x^T = [n, n_A, T]_{dev}$ and $u^T = [\dot{Q}, F_{out}]_{dev}$ are deviations from the indicated steady state. The terminal penalty $W(x)$ is the quadratic solution to an algebraic Ricatti equation for the linearized (13), and the local control law $\kappa(x)$ was derived from the optimal linear controller using the method presented in (DeHaan and Guay, 2005).

Using simple piecewise constant parameterizations, two sets of parameters ω^1 and ω^2 (both with $N = 15$, $\delta(x) \equiv 1s$, shown in Figure 2) were adapted online. The initializations ω_0^1 and ω_0^2 corresponded to different paths around the constraint region in the P-T plane (whose image is a skewed infinite cylinder in the x-space). As can be seen in Figure 1, neither initialization was very close to the infinite-horizon optimal solution. The controller initially selected the active parameterization ω^1 (i.e. "over" the constraint in Fig. 1), but as the dynamics and adaptation progressed, the active parameterization switched to ω^2 in time to feasibly pass under the constraint region.

6. CONCLUSIONS

Using existing results concerning discontinuous feedbacks, it has been demonstrated that a naive implementation of global search methods creates a robustness concern with respect to measurement noise for continuous-time (i.e. fast-sampled) MPC. While purely local realtime methods do not exhibit this characteristic, any mechanism for evaluating alternative paths generates a potential robustness concern. By incorporating hysteresis in the decision making, a framework has been created for improving the performance of realtime approaches in a nominally robust manner.

REFERENCES

Cannon, M. and B. Kouvaritakis (2000). Continuous-time predictive control of constrained nonlinear systems. In: *Nonlinear Model Predictive Control: Assessment and Future Directions for Research* (F. Allgower and A. Zheng, Eds.). Birkhauser-Verlag.

Chen, H. and F. Allgower (1998). A computationally attractive nonlinear predictive control scheme with guaranteed stability for stable systems. *Journal of Process Control* **8**(5-6), 475–485.

Clarke, F.H., Y.S. Ledyaev, R.J. Stern and P.R. Wolenski (1998). *Nonsmooth Analysis and Control Theory*. Grad. Texts in Math. 178, Springer-Verlag. New York.

Coron, J.M. and L. Rosier (1994). A relation between continuous time-varying and discontinuous feedback stabilization. *Journal of Mathematical Systems, Estimation, and Control* **4**(1), 67–84.

DeHaan, D. and M. Guay (2005). A new real-time method for nonlinear model predictive control. In: *International Workshop on Assessment and Future Directions of Nonlinear Model Predictive Control*. Freudenstadt-Lauterbad, Germany. pp. 73–84.

Grimm, G., M. Messina, S. Tuna and A. Teel (2003). Nominally robust model predictive control with state constraints. In: *Proc. IEEE Conf. on Decision and Control*. pp. 1413–1418.

Grimm, G., M. Messina, S. Tuna and A. Teel (2004). Examples when model predictive control is non-robust. *IEEE Trans. Automat. Contr.* **40**(10), 1729–1738.

Magni, L. and R. Sepulchre (1997). Stability margins of nonlinear receding-horizon control via inverse optimality. *Systems and Control Letters* **32**, 241–245.

Mayne, D. Q., J. B. Rawlings, C. V. Rao and P. O. M. Scokaert (2000). Constrained model predictive control: Stability and optimality. *Automatica* **36**, 789–814.

Messina, M., S. Tuna and A. Teel (2005). Discrete-time certainty equivalence output feedback: allowing discontinuous control laws including those from model predictive control. *Automatica* **41**(4), 617–628.

Ohtsuka, T. (2004). A continuation/GMRES method for fast computation of nonlinear receding horizon control. *Automatica* **4**(40), 563–574.

Tuna, S., R. Sanfelice, M. Messina and A. Teel (2005). Hybrid MPC: Open-minded but not easily swayed. In: *International Workshop on Assessment and Future Directions of Nonlinear Model Predictive Control*. Freudenstadt-Lauterbad, Germany. pp. 169–180.