PREDICTIVE CONTROL WITH ADDED FEEDBACK FOR FAST NONLINEAR SYSTEMS

E. Ronco, B. Srinivasan, J. Y. Favez, and D. Bonvin

Institut d'Automatique, Ecole Polytechnique Fédérale de Lausanne,

CH-1015 Lausanne, Switzerland

Email: bala.srinivasan@epfl.ch, Fax: +41 21 693 25 74

Keywords: Nonlinear systems; Predictive control; Optimal control; Neighbouring extremals.

Abstract

A major difficulty arises in predictive control when the prediction of the system evolution becomes unreliable due to the presence of uncertainty. This problem is particularly important for fast nonlinear systems since significant divergence can occur during the time required for optimization. As a solution to this problem, the addition of a linear state feedback controller regulating the prediction error is proposed. Using the neighbouring extremal theory, it is shown that this combination is a first-order approximation to an infinitely-fast receding horizon controller. The effectiveness of this methodology is illustrated *via* simulation on an inverted pendulum.

1 Introduction

Predictive control is a very effective approach for tackling problems with constraints and nonlinear dynamics, especially when analytical computation of the control law is difficult [5]. Standard predictive control involves recalculating at every sampling instant the input that minimises a criterion defined over a horizon window in the future, taking into account the current state of the system. Only the first part of the computed optimal input is applied to the system.

Such a methodology is widely used in the process industry where system dynamics are sufficiently "slow" to permit its implementation [7]. In contrast, applications of predictive control to fast dynamic systems are rather limited [8]. This is mainly due to the two following reasons: (i) the computational burden of the nonlinear optimization limits the frequency of reoptimization, and (ii) when the reoptimization interval is large relative to the system dynamics, the predictions are quite sensitive to modeling errors and disturbances.

Nevertheless, for small deviations from the optimal solution, a linear approximation of the system and a quadratic approximation of the cost are quite reasonable. In such a case, the theory of neighbouring extremals [2] provides a closed-form solution to the optimisation problem. Thus, the optimal input can be obtained using state feedback without explicit numerical optimisation. In this paper, a combination of predictive control and linear state feedback based on neighbouring extremals is proposed for the control of fast nonlinear systems.

Though the idea of using the theory of neighbouring extremals along with predictive control is novel, various studies have been reported in the literature that combine linear feedback and predictive control:

- *Dual mode predictive control [6]:* Predictive control is used when the state of the system is far away from the reference, while a linear feedback is used in its neighbourhood. Linear feedback guarantees the existence of a stabilizing input for times larger than the prediction horizon and thus helps in the stability analysis.

- *Feedback predictive control [5]:* In the presence of uncertainty, it is important for the purpose of robust optimization to be able to predict the dispersion of the states in the future caused by the uncertainty. This dispersion can be fairly large if the system is unstable. So, an additional feedback has been proposed to reduce this dispersion.

In dual mode predictive control, there is a switching between the predictive controller and the linear feedback controller, while both of them function together in feedback predictive control. With either approach, two cases can be distinguished depending on whether (i) this additional feedback is actually implemented [3], or (ii) a fictitious feedback, which is only used for the computation of the dispersion of states in the future, is considered [1].

In most methods available in the literature, the feedback controller is driven by the deviaton of the system response from a reference signal. In contrast, the proposed method

uses the deviation from a predicted state to drive the feedback controller. Since, due to model uncertainties, the model prediction does not quite correspond to the true evolution of the system, the system is forced to stay close to the predicted state. It is interesting to note the twist in the concept – the model is not adapted to provide a good prediction of the system, but instead, the input is adjusted so that the system follows the model prediction.

The paper is organized as follows. Background material regarding predictive control is presented in Section 2. The predictive controller enhanced with an additional feedback law is described in Section 3. Section 4 discusses some properties of this scheme, while Section 5 illustrates the advantages of the proposed predictive control scheme on a simulated inverted pendulum. Some concluding remarks are given in Section 6.

2 Preliminaries - Predictive Control

Consider the nonlinear system represented by:

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

where the state x and the input u are vectors of dimension n and m, respectively. x_0 represents the initial conditions and the function F describes the dynamics.

In predictive control, the following optimisation problem is solved repeatedly at discrete time t_k:

$$\min_{u([t_k, t_k + T_f])} J(t_k) = \frac{1}{2} x(t_k + T_f)^T P\, x(t_k + T_f) \qquad (2)$$

$$+ \frac{1}{2} \int_{t_k}^{t_k + T_f} \left(x(\tau)^T Q\, x(\tau) + u(\tau)^T R\, u(\tau) \right) d\tau$$

$$\dot{x} = F(x, u), \qquad x(t_k) = x_k$$

where P, Q, and R are positive-definite weighting matrices of appropriate dimensions, T_f is the prediction horizon, and x_k the states measured at time instant t_k. The optimal input obtained by solving (2) is represented by $u^\star([t_k, t_k + T_f])$.

Let δ be the optimisation time interval which, in general, is fixed in advance. The first part of the optimal input, $u^\star([t_k, t_k + \delta])$, is applied open-loop until the result from the next optimisation becomes available. Such an approach is referred to as intermittent predictive control [8]. For implementation purposes, the infinite dimensional input $u([t_k, t_k + T_f])$ needs to be parameterized using a finite number of decision variables. Furthermore, the optimiser can be constructed so as to account for the time delay δ in the computation of $u^\star([t_k, t_k + T_f])$.

Using Pontryagin's Minimum Principle (PMP), the problem of optimizing the *scalar* cost functional J in (2) can

be reformulated as the optimization of the Hamiltonian *function* $H(t)$ as follows [2]:

$$\min_{u(t)} H = \frac{1}{2}(x^T Q\, x + u^T R\, u) + \lambda^T F(x, u) \quad (3)$$

$$s.t. \quad \dot{x} = F(x, u), \quad x(0) = x_0 \qquad (4)$$

$$\dot{\lambda}^T = -H_x, \quad \lambda(t_f) = P\, x(t_f) \qquad (5)$$

where $\lambda(t) \neq 0$ is the n-vector of adjoint states (Lagrange multipliers for the system equations). The notation $a_b = \frac{\partial a}{\partial b}$ is used. The necessary condition for optimality is given by:

$$H_u = u^T R + \lambda^T F_u = 0 \qquad (6)$$

3 Predictive Control with Additional Feedback

When modeling errors and disturbances do not allow open-loop operation over the optimization interval δ, the input needs to be updated by means other than optimization. The measurements taken while the optimisation is in progress can be processed by a state feedback controller that is added to the predictive control scheme as shown in Figure 1. This results in control action generated by two feedback paths running at different speeds [9, 10].

- In the slower loop (encapsulated within the dotted box), the state of the system is used by the "Optimiser" every δ time units so as to compute the optimal input $u^\star([t_k, t_k + T_f])$, from which the nominal optimal input $u^\star(t)$ and the corresponding state evolution $x^\star(t)$ are generated using the "Trajectory generator".

- The faster loop consists of a linear state feedback controller that regulates $x(t)$ around $x^\star(t)$.
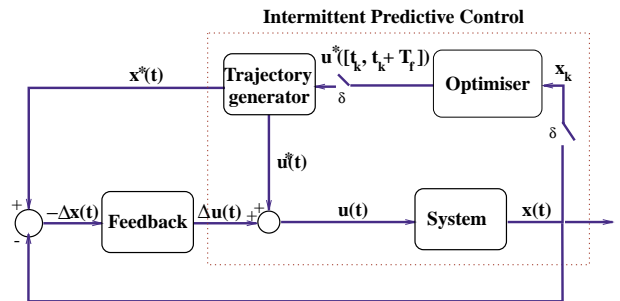
Figure 1: Predictive control with additional feedback

The control $u(t)$ applied to the system is the sum of $u^*(t)$ obtained from the optimizer and $\Delta u(t)$ computed through the linear feedback as follows:

$$\Delta u(t) = -K(t)\Delta x(t). \tag{7}$$

where $\Delta x(t) = x(t) - x^\star(t)$ and $\Delta u(t) = u(t) - u^\star(t)$.

If there is no uncertainty (modeling errors and disturbances), then $\Delta x(t) = 0$, and the linear feedback controller has no effect, $\Delta u(t) = 0$. The role of the feedback controller is therefore to steer the system towards the predicted state, thereby rejecting the effect of model uncertainty and disturbances. Thus, the input is adjusted so that the system follows whatever is predicted by the model, rather than the model being adapted to provide a good prediction for the system.

4 Analysis of the Proposed Scheme

This section examines the properties of predictive control enhanced with linear state feedback. Though, in general, there is no restriction on the type of feedback controller that could be used to regulate the prediction error $(-\Delta x)$, this study concentrates on linear state feedback controllers obtained from system linearisation.

4.1 Equivalence property

It will be shown using the theory of neighbouring extremals that the additional feedback performs optimisation implicitly, thereby acting as a standby whenever the numerical optimizer is unable to compute the optimal input sufficiently fast. The result is stated in the following theorem.

Theorem 1 *(Neighbouring extremal solution):*
For system (1) and the predictive controller (2), let $\bar{u}^\star(t)$ be the solution that would be obtained by repeating optimisation (2) continuously, while $u^\star(t)$ and $x^\star(t)$ are obtained by solving (2) once every δ time units. Let $\Delta x(t) = x(t) - x^\star(t)$ and $\Delta u(t) = u(t) - u^\star(t)$ be deviations from the computed optimal trajectory. If,

- *T_f is arbitrarily large,*
- *$x(t)$ is in the neighbourhood $x^\star(t)$, and*
- *the gain matrix $K(t)$ is solution to the matrix Ricatti equation*

$$\dot{S} = -H_{xx} + H_{xu}H_{uu}^{-1}H_{ux} + SF_uH_{uu}^{-1}F_uS -$$
$$S\left(F_x - F_uH_{uu}^{-1}H_{ux}\right) - \left(F_x - F_uH_{uu}^{-1}H_{ux}\right)^T S$$
$$S(t_k + T_f) = P \tag{8}$$

$$K = H_{uu}^{-1}\left(H_{ux} + F_u^T S\right) \tag{9}$$

then, to a first-order approximation,

$$\bar{u}^\star(t) = u^\star(t) - K(t)\Delta x(t) \tag{10}$$

Proof

Consider the time instant, $\tilde{t}$, in between two explicit optimizations, i.e., $\tilde{t} \in [t_k, t_{k+1}]$. First, it will be shown that the optimal input obtained from explicit optimization performed at $\tilde{t}$ can be approximated by the neighbouring extremal solution (10). Then, the argument will be extended to the case where $\tilde{t}$ takes all values in $[t_k, t_{k+1}]$.

The following optimization problem is solved at time instant $\tilde{t}$ to compute $\bar{u}^\star(\tilde{t})$:

$$\min_{u([\tilde{t},\tilde{t}+T_f])} J(\tilde{t}) = \frac{1}{2}x(\tilde{t}+T_f)^T P\, x(\tilde{t}+T_f) \tag{11}$$

$$+\frac{1}{2}\int_{\tilde{t}}^{\tilde{t}+T_f} \left(x(\tau)^T Q\, x(\tau) + u(\tau)^T R\, u(\tau)\right) d\tau$$

$$\dot{x}(t) = F(x,u) \qquad x(\tilde{t}) = x_{\tilde{t}}$$

where $x_{\tilde{t}}$ is to the measured state at time instant $\tilde{t}$. The criterion augmented with the constraints (system equations) reads:

$$\bar{J}(x,u) = \frac{1}{2}x(\tilde{t}+T_f)^T P\, x(\tilde{t}+T_f) + \tag{12}$$
$$\lambda(\tilde{t})^T x(\tilde{t}) - \lambda^T(\tilde{t}+T_f)x(\tilde{t}+T_f)$$
$$+\int_{\tilde{t}}^{\tilde{t}+T_f} \left(H(x,\lambda,u) + \dot{\lambda}^T x\right) d\tau$$

where the Hamiltonian function H is defined as in (3). Note that in the derivation of (12), $\int_{\tilde{t}}^{\tilde{t}+T_f} \dot{\lambda}^T x + \lambda^T \dot{x}\, d\tau = \lambda^T(\tilde{t}+T_f)\, x(\tilde{t}+T_f) - \lambda^T(\tilde{t})\, x(\tilde{t})$ is used.

Introducing the deviation variables $\Delta x(\tilde{t})$ and $\Delta u(\tilde{t})$ around $x^\star(\tilde{t})$ and $u^\star(\tilde{t})$, and under the assumption that the deviations are "small", the perturbation of the cost and the states from their optimal values can be captured using Taylor series expansions. Following the development in [2], the second-order expansion for the cost gives:

$$\bar{J}(x+\Delta x, u+\Delta u) = \bar{J}(x,u) \tag{13}$$
$$+ (Px(\tilde{t}+T_f) - \lambda(\tilde{t}+T_f))^T \Delta x(\tilde{t}+T_f) +$$
$$\lambda(\tilde{t})^T \Delta x(\tilde{t}) + \int_{\tilde{t}}^{\tilde{t}+T_f} \left((H_x + \dot{\lambda}^T)\Delta x + H_u\Delta u\right) d\tau$$
$$+ \frac{1}{2}\Delta x(\tilde{t}+T_f)^T P\, \Delta x(\tilde{t}+T_f) +$$
$$\frac{1}{2}\int_{\tilde{t}}^{\tilde{t}+T_f} [\Delta x^T\ \Delta u^T] \begin{bmatrix} H_{xx} & H_{xu} \\ H_{ux} & H_{uu} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix} d\tau$$

Note that the first-order variations are zero, since $H_u = 0$, $\dot{\lambda}^T = -H_x$, and $\lambda(\tilde{t} + T_f) = Px(\tilde{t} + T_f)$. In addition, the first-order expansion for the constraints of the optimisation problem (system equations) are considered. Also, the term $\lambda(\tilde{t})^T \Delta x(\tilde{t})$ is not affected by the input $u([\tilde{t}, \tilde{t} + T_f])$ and so is not considered for optimization. Then, the optimization problem (11) can be approximated in the deviation variables as:

$$\min_{\Delta u([\tilde{t}, \tilde{t}+T_f])} \frac{1}{2} \Delta x(\tilde{t} + T_f)^T P \, \Delta x(\tilde{t} + T_f) + \qquad (14)$$

$$\frac{1}{2} \int_{\tilde{t}}^{\tilde{t}+T_f} \begin{bmatrix} \Delta x^T & \Delta u^T \end{bmatrix} \begin{bmatrix} H_{xx} & H_{xu} \\ H_{ux} & H_{uu} \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta u \end{bmatrix} d\tau$$

$$\Delta \dot{x} = F_x \Delta x + F_u \Delta u$$

Problem (14) is the standard LQR problem whose optimal solution is given by $\Delta u_{opt}(t) = -K(t)\Delta x(t)$, for $t \in [\tilde{t}, \tilde{t} + T_f]$, where the gain of the feedback controller $K(t)$ is obtained from the solution of the Ricatti equation (8)-(9) but with $S(\tilde{t} + T_f) = P$ [2]. However, by the assumption that T_f is arbitrarily large, S has reached its steady state value and is constant over the interval $[t_k, t_k + \delta]$. So, it is sufficient to solve the Ricatti equation at t_k as in (8) and use it until the next explicit optimization is performed.

Thus, to a first-order approximation, the optimal solution to (11) $\bar{u}^\star(\tilde{t}) = u^\star(\tilde{t}) + \Delta u_{opt}(\tilde{t}) = u^\star(\tilde{t}) - K(\tilde{t})\Delta x(\tilde{t})$. The proof follows by extending the argument for all values of $\tilde{t} \in [t_k, t_{k+1}]$. ∎

Though, in this methodology, a simple linear feedback supplements predictive control, periodical recomputation of $u^\star(t)$ is indispensable. This realigns the computed optimal solution so that the actual state of the system is in its neighbourhood.

4.2 Stability

The stability of this cascade scheme is an important issue that needs attention. Theorem 1 states that the proposed scheme is, to a first-order approximation, equivalent to a single predictive control loop that operates infinitely fast (or at the sampling frequency of the fast loop). So, if the matrices P, Q, and R and the prediction horizon T_f are chosen so that this equivalent predictive controller stabilizes the system, then the proposed scheme will also stabilize the system. Alternatively, if it can be guaranteed that the cost decreases with every optimization, its cost function J can itself be used as a Lyapunov function to prove stability [5].

Though the aforementioned argument can be accepted from an intuitive point of view, no rigourous result can be stated for the moment due the approximations involved. The effect of these approximations on closed-loop stability forms one of the directions of future research.

4.3 Robustness

Though the methodology is motivated by the implementation of predictive control for fast systems, there is considerable interest to apply it to all systems, especially when the problem of robust stabilization is considered.

For robust stabilization, a terminal constraint of the form $x(t_k + T_f) \in \mathcal{X}$ is imposed to the optimization problem (2), and the input is computed in such a manner that the final state is in $\mathcal{X}$ despite uncertainty. The standard approach is to compute the dispersion of the states caused by the uncertainty using a model linearized around the nominal trjectory of the system and searching over open-loop input sequences that ensure that the family of trajectories end within $\mathcal{X}$. This is clearly overconservative since open-loop prediction is performed, while the true system evolution will take advantage of the closed-loop nature of predictive control [4]. When this linearized system is unstable, or has modes that do not vanish quickly, the system evolution and the dispersion of states grow with time. Then, for a large T_f, it becomes impossible to fit the set of all possible final states within $\mathcal{X}$. Also, it is not possible to choose a small T_f since a short prediction horizon can lead to closed-loop instability.

As mentioned in the introduction, two approaches are found in the literature to alleviate this problem: (i) explicit feedback (in addition to that inherently provided by predictive control) is used to reduce the dispersion of states [3], or (ii) fictitious feedback is used to estimate the dispersion of states to be expected with the predictive controller [1].

The neighbouring extremal approach presented in this paper fits both approaches. On the one hand, the linear state feedback computed using the neighbouring extremal theory helps reduce the dispersion of the future states. Thus, it can be used as the explicit feedback mentioned above. On the other hand, Theorem 1 indicates that performing optimization infinitely often is equivalent to using the neighbouring extremal control. Thus, since this linear feedback represents the inherent feedback provided by predictive control, it can be considered as the fictitious feedback that is used for computing a fair dispersion of the states. Note that the actual implementation is via repeated optimization only.

This new insight shows promise in the field of predictive control under uncertainty and constraints and needs further investigation.

5 Example: Control of an Inverted Pendulum

This section illustrates the properties of the predictive controller with additional feedback via the simulated control of an inverted pendulum.

5.1 System model

An inverted pendulum without the cart dynamics is considered in this section. The model equations are:

$$\dot{x_1} = x_2 \tag{15}$$

$$\dot{x_2} = \frac{m\,l}{J}[sin(x_1)\,g - cos(x_1)u] \tag{16}$$

where x_1 is the pendulum angle, x_2 its rotational velocity and u the control torque. The control objective consists of regulating the pendulum around the upright position starting from the downward position $x_0 = \begin{bmatrix} \pi & 0 \end{bmatrix}$. The following numerical values are used: $m = 1$, $g = 9.81$, $l = 1$ and $J = 1$. In addition, the control is constrained, $-50 < u < 50$. The sampling interval is 0.01s and the time taken for the numerical optimization is $\delta = 0.2$ s.

5.2 Control parameterization

This study follows [8] in that the input $u(t)$ is parametrised using

$$u(t) = \theta(t_k)\,e^{-\frac{(t-t_k)}{p}}, \tag{17}$$

where the time constant $p = 0.2$ s is chosen. The cost function considered has no integral cost, i.e., $Q = 0$ and $R = 0$. The prediction horizon is $T_f = 1$. A symbolic Gauss-Newton optimisation technique is used to minimise $J(t_k)$. Details of this approach can be found in [8].

5.3 Case of no uncertainty

In the absence of model errors and disturbances, control is straightforward, and the cost function quickly reduces to 0 with time (see Figure 2). Clearly, the additional feedback is inactive due to the absence of uncertainty, $\Delta x = 0$.

5.4 Case of uncertainty and no additional feedback

Uncertainty is introduced through both model mismatch and multiplicative measurement noise with standard deviation 0.05. Model mismatch is obtained by choosing
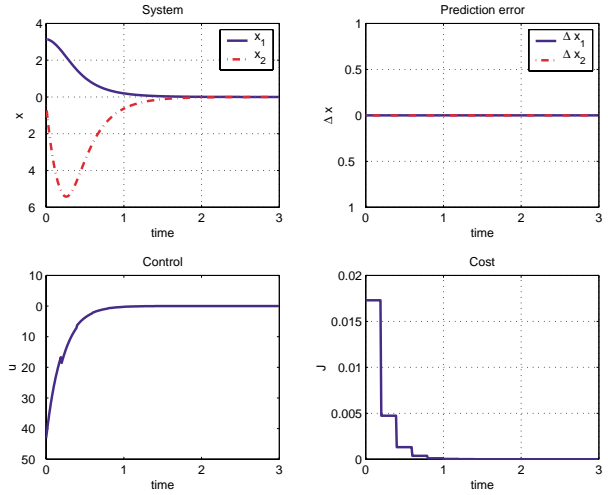

Figure 2: System responses, prediction error, control input, and cost in the case of no uncertainty

$\frac{m\,l}{g} = 0.7$ for the simulated reality. This makes the "system" much faster than the model.

The results using predictive control without additional feedback are depicted in Figure 3. The system does not settle down, and the cost function oscillates with time showing potential instability. This shows how a predictive controller can be destabilizing in the presence of uncertainty when the disturbances cannot be compensated sufficiently fast.
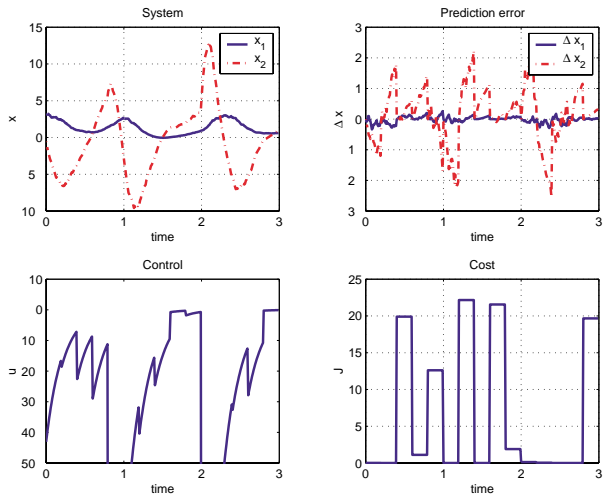

Figure 3: System response, prediction error, control input, and cost in the case of uncertainty and no additional feedback

5.5 Case of uncertainty and additional feedback

Linear time-varying feedback, which keeps the system in the neighbourhood of the predicted output, is added to the predictice controller as suggested in Section 4. The results are illustrated in Figure 4. Though the control is rather aggressive, the system settles down quickly. The system response is comparable to the one obtained in the absence of uncertainty (see Figure 2). This is the goal of the additional feedback, i.e., the actual system response is close to that predictied using the model, though the model is far from the reality. The cost quickly reduces to zero, implying stabilisation of the process. The prediction errors are much smaller than those obtained without additional feedback (see Figure 3). The additional feedback clearly increases the robustness of the predictive controller.
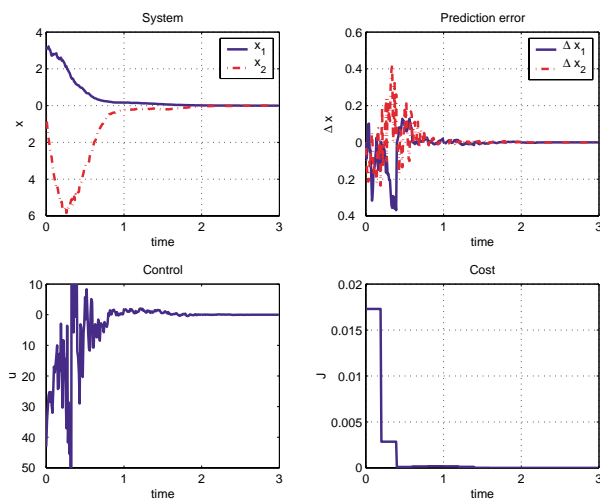

Figure 4: System response, prediction error, control input, and cost in the case of uncertainty and additional feedback

6 Conclusion

A method that helps reduce the effect of uncertainty in predictive control schemes dealing with fast nonlinear systems has been presented. During the time required for the computation of a new sequence of optimal inputs, a time-varying linear feedback controller regulates the prediction error using available measurements. Using the neighbouring extremal theory, it was shown that this combination is indeed a first-order approximation to an infinitely-fast receding horizon controller.

Predictive control with the addition of a fast feedback loop was found particularly effective to control the swing up of an inverted pendulum. The performance obtained with the proposed scheme in the presence of uncertainty is comparable to that obtained with standard predictive control in the absence of uncertainty.

Though the methodology needs further investigation, the results obtained so far are encouraging. Real time application of the proposed control startegy is envisaged.

References

[1] A. Bemporad. Reducing conservatism in predictive control of constrained systems with disturbances. In *37th IEEE Control and Decision Conference*, pages 1384–89, Tampa, FL, 1998.

[2] A. E. Bryson. *Dynamic Optimization*. Addison-Wesley, Menlo Park, California, 1999.

[3] B. Kovaritakis, J. A. Rossiter, and J. Schuurmans. Efficient robust predictive control. *IEEE Trans. Automat. Contr.*, 45(8):1545–49, 2000.

[4] J. H. Lee and Z. Yu. Worst-case formulations of model-predictive control for systems with bounded parameters. *Automatica*, 33(5):763–781, 1997.

[5] D. Q. Mayne, J. B. Rawlings, C. V Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[6] H. Michalska and D. Q. Mayne. Robust receding horizon control of constrained nonlinear systems. *IEEE Trans. Automat. Contr.*, 38(11):1623–33, 1993.

[7] S. J. Qin and T. A. Badgwell. An overview of industrial model predictive technology. In *Chemical Process Control V*, pages 232–256, Tahoe City, CA, 1997.

[8] E. Ronco, T. Arsan, and P. J. Gawthrop. Open-loop intermittent feedback control: Practical continuous-time gpc. *IEE Proc. on Control Theory and Applications*, 146(5):426–434, 1999.

[9] P. Terwiesch and M. Agarwal. On-line correction of pre-determined input profiles for batch reactors. *Comp. Chem. Eng.*, 18:S433–S437, 1994.

[10] E. Visser, B. Srinivasan, S. Palanki, and D. Bonvin. A feedback-based implementation scheme for batch process optimization. *J. Process Contr.*, 10:399–410, 2000.