

Interaction between control and state estimation in nonlinear MPC

MORTEN HOVD* and ROBERT R. BITMEAD†

This paper shows how to take the quality of the state estimation into account in the formulation of the optimization criterion for model predictive control (MPC). This is relevant for the control of nonlinear plants, for which the separation principle in general does not apply.

1. Introduction

Model predictive control (MPC) has found wide-spread application, particularly in the chemicals processing industry and refineries. Originally, MPC was proposed using linear plant models, with constraints on inputs and states (or outputs) being the only type of nonlinearity that was addressed. However, in the last two decades, there has been a number of publications on MPC using nonlinear process models, Allgöwer *et al.* (1999) and a number of actual industrial applications using such nonlinear models, Qin & Badgwell (2003).

The typical academic publication on nonlinear MPC focuses on developing stability guarantees in the face of plant nonlinearities, based on the assumption that the plant states are available either through direct measurements or through the use of a state estimator.

In the processing industries, which has seen the bulk of MPC applications so far, one seldom enjoys the luxury of being able to measure all states directly. One therefore has to resort to some sort of state estimation, the Extended Kalman Filter (EKF) or Receding Horizon Estimation (RHE) being the estimators most commonly proposed.

Although the control methodology (nonlinear MPC) and state estimation methodology (EKF/RHE) are powerful, this approach totally ignores the fact that the separation principle in general does not apply to nonlinear plants. In Imsland *et al.* (2003) and Adetola & Guay (2003) it is shown how the performance of an output feedback NMPC controller can achieve the performance and region of attraction of a state feedback NMPC controller. However, these results are restricted to highly specific nonlinear model structures and requires very high observer gains.

This paper shows how to take into account the quality of the state estimate in the formulation of the optimization problem for the MPC. In previous work, Yan and Bitmead (2004) have shown how to include the quality of the state estimate into the constraint formulation for MPC. Here we will include the effects of manipulated variable moves on the state estimate covariance in both the constraints and the optimization criterion. This will of course result in control moves that differ from those

*Department of Engineering Cybernetics, Norwegian University of Science and Technology, N-7491 Trondheim, Norway, e-mail: morten.hovd@itk.ntnu.no. The paper was presented at the IFAC conference DYCOPS7, Boston, Massachusetts, USA, July 2004.

†Department of Mechanical and Aerospace Engineering, University of California San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0411, USA.

that would be obtained if the estimation accuracy was ignored, i.e., it will 'degrade' the control to obtain more accurate state estimates. However, insisting on accurate control based on inaccurate state estimates is not very meaningful. The proposed modification is clearly only relevant for nonlinear plants, since for linear plants the quality of the state estimation is independent of the control action.

In section 2 the model predictive control formulation is presented, and the state estimation approaches (EKF and RHE) are presented briefly. The proposed control and estimation method is illustrated on a simple example in section 3. The properties of the proposed MPC formulation are addressed in section 4. The final section presents conclusions and points out directions for further work.

2. MPC and state estimation formulations

2.1. MPC formulation

The MPC formulation is based on a discrete nonlinear state space model of the type

$$x_{k+1} = f(x_k, u_k) + w_k \quad (1)$$

$$y_k = g(x_k, u_{k-1}) + v_k \quad (2)$$

where x is the state vector, u is the input vector, y is the measurement vector, whereas w and v are state excitation noise and measurement noise, respectively. The noise sequences are assumed to be Gaussian white noise. The subscript is a sampling time index. It is assumed that the measurement is obtained before calculating the new value for the input for a given sample interval, therefore y_k may depend on u_{k-1} but not on u_k .

The proposed MPC may be formulated as

$$\begin{aligned} & \min_{u_k^T, u_{k+1}^T, \dots, u_{k+n_p-1}^T} \quad (3) \\ & \sum_{i=1}^{n_p} (x_{k+i} - x_{k+i}^r)^T Q_{k+i} (x_{k+i} - x_{k+i}^r) \\ & + \sum_{i=1}^{n_p} (u_{k+i-1} - u_{k+i-1}^r)^T R_{k+i} (u_{k+i-1} - u_{k+i-1}^r) \\ & + \sum_{i=1}^{n_p} \sum_{m=1}^{n_v} \rho_{i,m} v_m^T P_{k+i} v_m \end{aligned}$$

such that

$$H_c x_{k+i} + G_c u_{k+i} + L_c \sigma_{c,k+i} \leq b_c$$

$$x_{k+i} = f(x_{k+i-1}, u_{k+i-1})$$

The matrix P_{k+i} is the *a posteriori* EKF state covariance matrix at time $k+i$. This will sometimes be denoted $P_{k+i|k+i}$ to clearly distinguish it from the *a priori* covariance matrix $P_{k+i|k+i-1}$. Although the EKF state covariance matrix is not an exact representation of the state probability distribution for nonlinear systems, it is a measure of state estimation accuracy that can be obtained with relative ease. Note also that P_{k+i} is easily predicted along a future state trajectory. The vectors v_m are typically chosen to be orthonormal, and allow using different weights on the state covariance for different directions of the state space. In the constraint equations, $\sigma_{c,k+i}$ represents the predicted standard deviation of the state estimate in the direction of relevance for constraint c . No

constraint can be guaranteed to hold in the future if the plant is subject to stochastic inputs (at least if these stochastic inputs are unbounded, as is the case for Gaussian noise). However, by including the standard deviation of the state estimate in the constraints, one has a handle for tuning the controller to ensure an acceptable probability for constraint violation.

The vectors x^r and u^r are reference vectors for the states and manipulated inputs, respectively. These reference vectors are assumed to be known *a priori*.

To enable efficient computation in on-line applications, there is a practical requirement that the optimization problem must be solvable within one sampling period. Therefore, MPC optimization problems are in general formulated as LP or QP problems. Problems with significant non-linearities are typically approximated by QP problems, with line search along the solution vector to improve both the accuracy and the rate of convergence to the solution, Biegler (1998). Further improvements in accuracy may be obtained by using the result of one QP with line search as the starting point for a new optimization, if computing power and available time allow.

In Eq. (3), it is the non-linear plant model which makes the optimization problem differ from a standard QP problem. For linear plant models (subject to Gaussian noise) the state estimate covariance matrix and corresponding state estimate standard deviation would be independent of the future input and state trajectories.

In order to approximate Eq. (3) with a QP problem, we will linearize the plant model around given future state and input trajectories, denoted by x_{k+i}^0 and u_{k+i}^0 , respectively. Typically, these future state and input trajectories are obtained from by simulating of the non-linear plant model n_p timesteps into the future, starting from the current state estimate x_k , and applying the 'tail' of the vector of manipulated variables obtained at the previous time step (assuming zero state excitation noise w_{k+i}). In addition to providing future trajectories for the states and manipulated variables, such a simulation can also provide nominal future values for the EKF state covariance matrix P_{k+i}^0 . Thus, the nonlinear plant model, Eqs. (1, 2), are replaced in the approximate QP problem by the time varying linear model

$$x_{d,k+i+1} = A_{k+i}x_{d,k+i} + B_{k+i}u_{d,k+i} \quad (4)$$

$$y_{d,k+i} = C_{k+i}x_{d,k+i} + D_{k+i}u_{d,k+i-1} \quad (5)$$

The subscript d on the measurements, states, and manipulated variables indicate that these are here the *deviations* from the values around which the linearization is performed.

Having thus 'removed' the non-linearity in the constraints that are caused directly by the nonlinear plant model, we must next do the same for the predicted values of the state covariance matrix in the optimization criterion, and the state standard deviation in the constraints. This is done by a Taylor series expansion around the nominal future state and manipulated variable trajectories. We will henceforth adopt the more economic notation

$$\begin{aligned} \chi &= [x_{d,k+1}^T, \dots, x_{d,k+n_p}^T]^T \\ v &= [u_{d,k}^T, \dots, u_{d,k+n_p-1}^T]^T \\ \chi_0 &= [(x_{k+1}^0)^T, \dots, (x_{k+n_p}^0)^T]^T \\ v_0 &= [(u_k^0)^T, \dots, (u_{k+n_p-1}^0)^T]^T \end{aligned}$$

Thus, we get

$$\begin{aligned} v_m^T P_{k+i} v_m &\approx v_m^T P_{k+i}^0 v_m + dP_{\chi, m, k+i} \chi \\ &+ dP_{v, m, k+i} v + \frac{1}{2} [\chi^T v^T] ddP_{m, k+i} [\chi] \end{aligned} \quad (6)$$

where

$$dP_{\chi, m, k+i} = \frac{\partial(v_m^T P_{k+i} v_m)}{\partial \chi}, \quad dP_{v, m, k+i} = \frac{\partial(v_m^T P_{k+i} v_m)}{\partial v}$$

and $ddP_{m, k+i}$ is the corresponding matrix of second partial derivatives

$$ddP_{m, k+i} = \begin{bmatrix} \frac{\partial^2(v_m^T P_{k+i} v_m)}{\partial \chi^2} & \frac{\partial^2(v_m^T P_{k+i} v_m)}{\partial \chi \partial v} \\ \frac{\partial^2(v_m^T P_{k+i} v_m)}{\partial v \partial \chi} & \frac{\partial^2(v_m^T P_{k+i} v_m)}{\partial v^2} \end{bmatrix} \quad (7)$$

The vectors v_m are independent of states and manipulated variables, and the evaluation of the partial derivative matrices defined above therefore requires the evaluation of the partial derivatives of P_{k+i} . To do this, we need to start from the familiar EKF covariance equations

$$\begin{aligned} P_{k+i|k+i-1} &= A_{k+i-1} P_{k+i-1|k+i-1} A_{k+i-1}^T \\ &+ W \end{aligned} \quad (8)$$

$$\begin{aligned} P_{k+i|k+i} &= P_{k+i|k+i-1} \\ &- P_{k+i|k+i-1} C_{k+i}^T M C_{k+i} P_{k+i|k+i-1} \end{aligned} \quad (9)$$

where $M = (V + C_{k+i} P_{k+i|k+i-1} C_{k+i}^T)^{-1}$. The next step is to differentiate the linearized model equations around the predicted future state and input trajectories. Note that the linearized model equations will only depend on the state and input values at the same timestep. The effects of states and inputs at previous timesteps are accounted for through the model equations themselves, whereas future states and inputs obviously cannot have any effect. On the other hand, the EKF state covariance matrix will depend on deviations from the nominal state and input trajectories at previous time steps, through the EKF covariance equations above.

Thus, we have

$$A_{k+i}^{x_j} = \left. \frac{\partial A_{k+i}}{\partial x_{j, k+i}} \right|_{x_{k+i}^0, u_{k+i}^0} \quad (10)$$

$$A_{k+i}^{u_j} = \left. \frac{\partial A_{k+i}}{\partial u_{j, k+i}} \right|_{x_{k+i}^0, u_{k+i}^0} \quad (11)$$

$$C_{k+i}^{x_j} = \left. \frac{\partial C_{k+i}}{\partial x_{j, k+i}} \right|_{x_{k+i}^0, u_{k+i}^0} \quad (12)$$

$$C_{k+i}^{u_j} = \left. \frac{\partial C_{k+i}}{\partial u_{j, k+i-1}} \right|_{x_{k+i}^0, u_{k+i-1}^0} \quad (13)$$

Here the index j refers to a particular element of the state or manipulated variable vector at time $k+i$. Note the time index on u in Eq. (13) above. This reflects the assumed order of execution of tasks at each sampling instant:

1. Obtain measurements.
2. Update state estimates.

3. Calculate new manipulated variables.

Thus, the manipulated variable that are applied at the time when y_{k+i} is obtained, is actually u_{k+i-1} .

Tedious, but straight forward application of matrix calculus will then produce the needed partial derivatives $dP_{x,m,k+i}$, $dP_{v,m,k+i}$, and $ddP_{m,k+i}$. Due to space limitations, this derivation is omitted from the paper.

What remains before the MPC problem can be formulated is to linearize the state standard deviation in the constraints in Eq. (3). Thus, we use the approximation

$$\sigma_{c,k+i} \approx \sigma_{c,k+i}^0 + \left[\frac{\partial \sigma_{c,k+i}}{\partial \chi} \frac{\partial \sigma_{c,k+i}}{\partial v} \right] \begin{bmatrix} \chi \\ v \end{bmatrix} \quad (14)$$

Let h'_c be row r of the matrix H_c in Eq. (3). The standard deviation of the states in the direction perpendicular to the constraint boundary is then

$$\sigma'_{c,k+i} = \sqrt{h'_c P_{k+i} (h'_c)^T} \quad (15)$$

Thus,

$$\begin{aligned} \left[\frac{\partial \sigma_{c,k+i}}{\partial \chi} \frac{\partial \sigma_{c,k+i}}{\partial v} \right] &= \frac{1}{\sqrt{h'_c P_{k+i} (h'_c)^T}} \times \\ &\left[\frac{\partial (h'_c P_{k+i} (h'_c)^T)}{\partial \chi} \frac{\partial (h'_c P_{k+i} (h'_c)^T)}{\partial v} \right] \end{aligned} \quad (16)$$

Hence, we will also here require partial derivatives of P_{k+i} w.r.t. states and manipulated variables.

Everything is now ready for formulating an approximation to Eq. (3) as a QP problem. The required manipulations are fairly standard, and are not repeated here. The MPC computations at each timestep are organized as follows:

1. The 'tail' of the manipulated variable vector from the previous timestep is used as an initial guess for the solution.
2. This manipulated variable sequence is applied to the nonlinear model in Eq. (1), and the model is linearized around the manipulated variable and the state trajectories.
3. The QP problem is formulated and solved.
4. A line search along the solution vector from the QP problem is performed. In the line search, the actual nonlinear model and nonlinear predictions of the EKF covariance matrices are used.

2.2. State estimation

Two related approaches are used for state estimation, the Extended Kalman Filter (EKF) and Receding Horizon Estimation (RHE).

For the EKF, standard formulae are used, using local linearizations of the system dynamics. The EKF can be found in many textbooks on advanced control.

For RHE, a formulation similar to the constrained linear Moving Horizon Estimator of Rao (e.g. Rao (2000)), is used. That is, a QP formulation is used to approximate the observed system outputs over the estimation horizon, with the linearized equations describing the system dynamics as constraints. The free variables in the optimization are measurement and state excitation noise sequences, and the state at the beginning of the

estimation horizon. The arrival cost (i.e., weight on the state at the beginning of the estimation horizon) is based on the EKF covariance, as described by Rao (2000).

The QP problem gives us a state excitation noise sequence, as well as the change from the *a priori* to the *a posteriori* state estimate at the start of the estimation horizon. For any multiple of these QP outputs, the system may be simulated using the nonlinear system model, and the corresponding measurement noise sequence required to obtain the observed measurements can be calculated. This is used in a line search procedure after solving the QP problem, in much the same way as line search is commonly used in nonlinear MPC.

3. Example

The method is applied to a simple example, given by

$$x_{k+1} = x_k + 0.1x_k u_k + u_k + w_k \quad (17)$$

$$y_k = x_k^3 + v_k \quad (18)$$

The state excitation noise w is assumed to be Gaussian with variance 0.01, whereas the measurement noise v has variance 0.1. The MPC has prediction horizon $n_p = 10$, state weight $Q_{k+i} = 2$ for $i = 1, \dots, 9$, $Q_{k+10} = 10$, and $R_{k+i} = 1$. All simulations are started from an initial state $x_0 = 0.5$ and a small initial state covariance $P_0 = 0.01$. Figure 1 shows the results when not taking the covariance of the state estimate into account ($\rho_i = 0$ in Eq. (3)). The state estimation is performed with an Extended Kalman Filter. We see that the controller quickly controls the estimated state to zero. Once the estimated state is zero, the system is no longer observable, and the EKF can no longer correct for the observed offset between actual and predicted measurement. Consequently, the state starts to drift. In Figure 2 we see the corresponding results when a value of $\rho_i = 0.3$ is used. The manipulated variable starts moving when the state covariance starts becoming large. This moves the state estimate into a region where the system is observable, and the EKF can correct the state. It may not be immediately obvious that the performance achieved in Figure 2 is better than that achieved in Figure 1. However, over a sufficiently long operating period, the actual state in Figure 1 may drift arbitrarily far from the reference, whereas this does not happen in Figure 2.

If the state estimate in Figure 2 has drifted any significant distance from the reference (at which the system is unobservable), the gain from state to measurement will be much larger at the actual state value than at the estimated state. The EKF, using a model linearized around the estimated state, therefore modifies the state estimate excessively. This is particularly evident just after 300 timesteps in Figure 2. Figure 3 shows the results obtained when using Receding Horizon Estimation with an estimation horizon of 10. The main difference between the performance of the EKF and the RHE is here that a line search along the solution vector is used in the RHE. Since the nonlinear model is used in the line search, the overcorrection that was evident in the EKF is significantly reduced. On the other hand, a somewhat more 'noisy' state estimate is obtained, leading to more active use of the manipulated variables. Nevertheless, Figure 3 clearly shows the best performance among the three cases shown here. Note that since the measurement function is the state cubed, the overshoot shown in Figure 2 would appear much more dramatic if the measurement and not the state were shown.




Figure 1. Results when using no weight on the EKF state covariance estimate.

4. Properties of the proposed method

In this paper, a method is proposed for taking the quality of the state estimate into account in the optimization problem in MPC. This consideration is only of relevance for nonlinear systems; It is well known that the control and estimation problems can be separated for linear systems. The theoretical properties of this type of MPC controller has not been extensively studied. However, some relevant issues are discussed below.

4.1. Asymptotic stability

A formal stability analysis for the type of MPC controller presented here is likely to be hard. Clearly, this type of controller will not result in an *asymptotically* stable closed loop system, if the system is not observable at the reference state x^* (and corresponding input u^*). This is not a shortcoming of the proposed control method, though, it merely means that the system needs to be excited (i.e., moved away from the reference state) in order to update the state estimate. For acceptable control to be achievable, it must be possible to obtain acceptable state estimates in a region not too far from the reference state.

4.2. Computational issues

4.2.1. Boundedness of solutions Obviously, it is a practical requirement that the solution to the optimization problem should be bounded. For convex QP problems this is automatically ensured by the shape of the objective function itself. For a non-convex




Figure 2. Results when using the weight $\rho = 0.3$ on the EKF state covariance estimate.

problem, this must be ensured by the constraints. Applying maximum and minimum value constraints to the manipulated variables will automatically ensure this. For most problems, such constraints are imposed by the physics of the problem anyway.

4.2.2. Required computation time Typical MPC formulations are formulated as LP or convex QP problems. For the method proposed here, it is possible that the problem can lose convexity if the EKF state covariance estimate grows so large as to dominate the cost of control error and manipulated variable moves. Non-convex QP problems are generally 'np-hard', meaning that the required time to guarantee that the optimal solution is found will in general grow exponentially with problem size. Thus, it may take unacceptably long to find the global optimum even for problems of moderate size. However, this does not imply that it will necessarily take long to find an acceptable solution. The constraints used in this MPC formulation are linear, and thus finding a *feasible* starting point is no harder than with traditional MPC formulations. Any gradient-based optimization code can then be used to find a local optimum—what is hard is to ascertain whether such a local optimum is also globally optimal. Even if a solution that is only locally optimal is used in the control, the resulting control performance may well be acceptable.

4.2.3. The need for line search If the weight on the EKF covariance matrix is set to zero, the objective function is by construction convex, i.e., it resembles a 'bowl' in shape. Consider a case where the plant is locally weakly unobservable at the reference state, like in the example above. Keeping the (estimated) state at the reference will then




Figure 3. Results when using the weight $\rho_i = 0.3$ on the EKF state covariance estimate in the MPC, while RHE is used for state estimation.

necessarily mean that the state covariance grows with time. With a non-zero weight on the covariance matrix, this will first flatten the bottom of the bowl described by the objective function, and then introduce a 'hill' around the reference state. Formulating the optimization problem as a QP problem means that we are locally fitting a quadratic function to approximate the true objective function. Fitting the quadratic function to the top of the 'hill' in the bottom of the 'bowl' can clearly lead to gross errors when the solution is based on the fitted quadratic function only. Performing a line search along the direction specified by the solution vector is the standard way to guard against this type of problems in nonlinear optimization.

5. Conclusions and Further Work

This work presents a novel way of including state estimate quality into the MPC formulation. It represents a significant extension of the work of Yan and Bitmead (2004). A topic for further research is to extend this approach to simultaneous state and parameter estimation, which would be a step towards a dual MPC controller, since it has the capability to generate probing signals to better identify states and parameters. This will also introduce back-off from constraints due to uncertain model parameters, while the controller will actively probe the control plant in order to identify the parameters. Note that in this extension of our work, the probing will only be active when parameter estimates are uncertain, and the probing will only excite directions affected by the uncertain parameters. This is in contrast to the Model Predictive Control and

Identification (MPCI) formulation proposed in, Shouche *et al.* (2002), which imposes a persistence of excitation constraint, and hence excites all model directions continuously.

The MPC formulation proposed herein does require heavy on-line computation. Further research is needed to analyze how the resulting optimization problem can be solved efficiently. Emerging results on explicit MPC formulations for nonlinear systems (see e.g. Johansen (2002)) appears to be a promising direction for further research. This approach may also provide more control on the degree of sub-optimality of the implemented controller.

References

- ADETOLA, V. & GUAY, M. (2003). Nonlinear output feedback receding horizon control of sampled data systems. In *Proc. of the American Control Conference*, pp. 4914–4919, Denver, Colorado, June.
- ALLGÖWER, F., BADGEWELL, T. A., QIN, J. S., RAWLINGS, J. B. & WRIGHT, S. J. (1999). Nonlinear model predictive control and moving horizon estimation—an introductory overview. In *Advances in Control. Highlights of the ECC'99*. Springer.
- BIEGLER, L. T. (1998). Advances in nonlinear programming concepts for process control. *Journal of Process Control*, 8, pp. 301–311.
- IMSLAND, L., FINDEISEN, R., BULLINGER, E., ALLGÖWER, F. & FOSS, B. A. (2003). A note on stability, robustness and performance of output feedback nonlinear model predictive control. *Journal of Process Control*, 13, pp. 633–644.
- JOHANSEN, T. A. (2002). On multi-parametric nonlinear programming and explicit nonlinear model predictive control. In *IEEE Conf. Decision and Control*, pp. 2768–2773, Las Vegas, NV, USA.
- QIN, S. J. & BADGEWELL, T. A. (2003). A survey of industrial model predictive control technology. *Control Engineering Practice*, pp. 733–764.
- RAO, C. V. (2000). *Moving Horizon Strategies for the Constrained Monitoring and Control of Nonlinear Discrete-Time Systems*. PhD thesis, Department of Chemical Engineering, University of Wisconsin–Madison.
- SHOUCHE, M. S., GENÇELİ, H. & NIKOLAOU, M. (2002). Effect of on-line optimization techniques on model predictive control and identification (MPCI). *Computers & Chemical Engineering*, 26, pp. 1241–1252.
- YAN, J. & BITMEAD, R. R. (2004). Incorporating state estimation into model predictive control and its application to network traffic control. *Accepted for publication in Automatica*.