The 4th Asian Control Conference
September 25-27, 2002
Singapore

A Comparison of Optimal Control Strategies for a Toy Helicopter

Jonas Balderud* and **David I. Wilson**†

Dept. of Electrical Engineering,
Karlstad University, Sweden
e-mail: *jonas.balderud@kau.se, †david.wilson@kau.se

Abstract

This work compares the performance of three optimal controllers on a toy helicopter: (i) model predictive control, (ii) linear quadratic optimal control combined with a state estimator, and (iii) optimal linear quadratic *output* control. These three schemes obtained significantly improved results over that achievable using classical control algorithms.

Keywords: Optimal output control, PLQ, LQG, model predictive control, helicopter

1 Introduction

The multivariable toy helicopter from Humusoft, CZ shown Fig. 1 makes for an impressive, visually arresting, control demonstration, [1]. Developing optimal controllers for the open-loop unstable plant, that offer convincingly superior performance to classical control algorithms when implemented on actual non-trivial hardware is something that does not receive the due attention it perhaps should.

This paper compares the application of three different types of popular optimal controllers on the helicopter. These are: (i) model predictive control (MPC) [2], (ii) linear quadratic optimal control combined with a state estimator, and (iii) optimal linear quadratic *output* control, [3, 4]. These three schemes were applied to an actual bench-scale helicopter to obtain significantly improved results over that achievable using classical control algorithms.

The material in the paper is organised as follows. Section 2 describes the helicopter, highlights the problems using classical control algorithms, and describes a semi-physical dynamic model of the helicopter suitable for control. Section 3 outlines the three optimal control schemes compared in this study: Model Predictive Control (MPC), Linear Quadratic Control with state estimation (LQG), and optimal output control. Section 4 compares the performance of the various optimal control scheme and section 5 finishes with some conclusions.

Figure 1: The bench-scale 2DOF helicopter with joystick.

2 The helicopter and classical control

The helicopter in Fig. 1 has two degrees of freedom, (elevation, azimuth), and three inputs (main and side rotors, and a moveable counter weight) and is intended for education in automatic control, [1]. A detailed explanation of one particular optimal control strategy has been previously reported in [5, 6]. This non-square multi-input/multi-output bench-scale apparatus is challenging to control because of the fast nonlinear dynamics, severe stiction in the bearings, the stochastic nature of the disturbances and the strong interactions in the plant which exhibits both stable and unstable modes. Fig. 2 illustrates the multiple equilibria in the elevation dynamics; below the horizontal the plant is stable, above, unstable.

Classical control such as two PID controllers suffer due to the strong nonlinearities evident from the openloop responses in Fig. 2, and from the obvious coupling between the elevation and azimuth dynamics. Fig. 3 shows a typical response when using two PID controllers at a sampling time of $\Delta t = 5$ms. The derivative term, necessary to stabilise the oscillations, is evident in the excessive movement of the controller input.

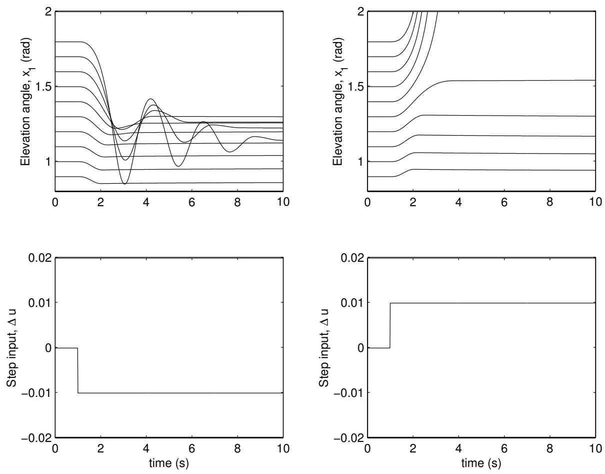

Figure 2: Various openloop step changes in elevation. Above the horizontal, the elevation dynamics are unstable.

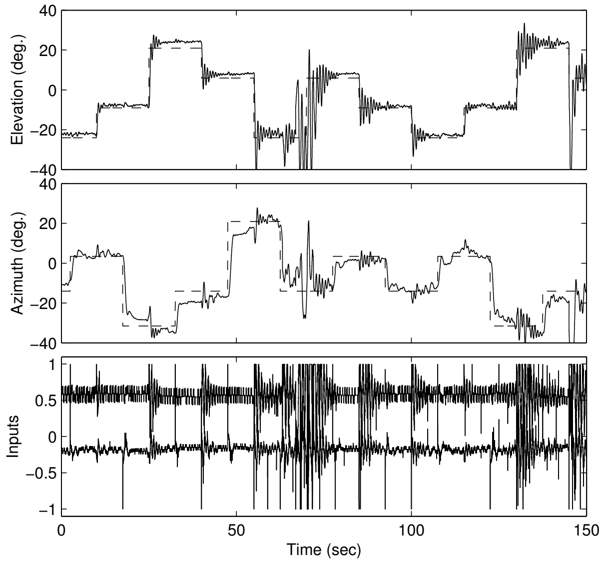

Figure 3: Closed-loop response of the helicopter using 2 PID controllers with a $\Delta t = 5$ms.

2.1 Helicopter model

The improved performance of optimal controllers are a direct consequence of the dynamic model. Four states are needed to model the position and velocity of the body in the two axes,$(\theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2)$, and two further states, (ω_1, ω_2), are needed to model the two motor dynamics giving a total of six, coupled, highly-nonlinear states. A semi-rigorous model incorporating friction,

$$\begin{aligned} I_m \ddot{\theta}_1 =& K_1\omega_1^2 - (\beta_{11}|\omega_1| + \beta_{21})\dot{\theta}_1 \\ & - T_{c_1}\mathrm{sign}(\dot{\theta}_1)\left(1 - \mathrm{e}^{-(|\dot{\theta}_1|/\dot{\theta}_{1_0})}\right) \\ & - T_g \sin(\theta_1 + \alpha_1) - K_G\dot{\theta}_2\omega_1\cos(\theta_1) \\ & + K_C\dot{\theta}_2^{\,2}\cos(\theta_1) \end{aligned} \tag{1}$$

$$I_{r_1}\dot{\omega}_1 = u_1 - a_1\omega_1^2 - b_1\omega_1 \tag{2}$$

$$\begin{aligned} \sin(\theta_1) I_{s_0}\ddot{\theta}_2 =& \sin(\theta_1)K_2|\omega_2|\omega_2 - (\beta_{12}|\omega_2| + \beta_{22})\dot{\theta}_2 \\ & - T_{c_2}\mathrm{sign}(\dot{\theta}_2)\left(1 - \mathrm{e}^{-(|\dot{\theta}_2|/\dot{\theta}_{2_0})}\right) \\ & - \sin(\theta_1 + \alpha_2) \\ & K_{r_1}\left(u_1 - K_{r_2}\left(a_1\omega_1^2 + b_1\omega_1\right)\right) \end{aligned} \tag{3}$$

$$I_{r_2}\dot{\omega}_2 = u_2 - a_2|\omega_2|\omega_2 - b_2\omega_2 \tag{4}$$

was developed in [6] which also details the identification of the 23 plant parameters when the counter weight is at 100% full scale. The two outputs are the axis rotations θ_1 and θ_2 and a Kalman filter was used to reconstruct $\dot{\theta}_1, \dot{\theta}_2, \omega_1, \omega_2$ needed for full state feedback. The 95% settling time for the elevation dynamics is around 3–6s indicating a suitable sampling time of 0.05s. In the following, the nominal plant is when the counter weight is at 100% full scale, while the model/plant mismatch case is when the counter weight is moved aft to 75% full scale.

3 Three optimal control schemes

3.1 Model predictive control

Model Predictive Control (MPC) refers to a family of controllers that use a model to compute an input trajectory in order to optimize the future behaviour of the plant, [2, 7, 8]. The optimization is repeated each sample, as only the first future manipulated variable adjustment is actually implemented on the plant. Linearising Equations (1)–(4) about the current operating point with a differenced input gives

$$\mathbf{x}(k+1) = \mathbf{A}\mathbf{x}(k) + \mathbf{B}\Delta\mathbf{u}(k) \tag{5}$$

$$\mathbf{y}(k) = \mathbf{C}\mathbf{x}(k) \tag{6}$$

The control law finds the next N_c input moves over a longer prediction horizon, N_p in order to minimise a weighted sum of output deviations from a desired trajectory, $\mathbf{r}$, and input trajectory. For a linear plant, the analytical solution

$$\Delta\mathbf{u} = \left(\mathbf{H}^T\mathbf{H} + \mathbf{\Lambda}\right)^{-1}\mathbf{H}^T(\mathbf{r} - \mathbf{y}_f) \tag{7}$$

gives the control moves as a function of future setpoints, the free response of the plant, $\mathbf{y}_f$ and the block Hankel matrix $\mathbf{H}$ which in turn is composed of the linear plant model matrices. As the current state appears implicitly in the term $\mathbf{y}_f$ in Equation (7), a state observer is needed to complete the control algorithm. The state observer, such as a Kalman Filter, provides direct feedback into the control law, which is different from the somewhat ad-hoc solution used by the DMC, [9], algorithm. Note that it was found necessary to integrate the full nonlinear model as proposed by [10] to compute the free response as integrating the linear model proved unstable.

3.1.1 MPC performance: Fig. 4 shows a typical servo response using an MPC controller with future knowledge of the setpoint changes (acausal) and with the counter weight in the nominal position. The choice of the prediction $N_p = 80$, and control $N_c = 30$ horizons were obtained from prior tests.

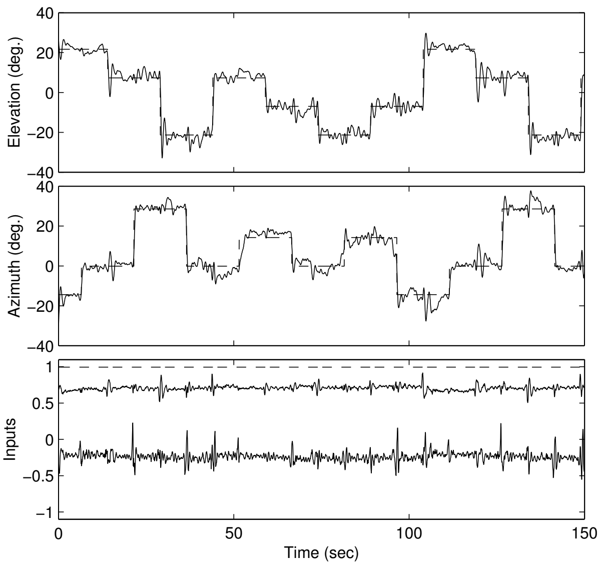

Figure 4: Acausal MPC controlled response with $N_p = 80, N_c = 30$. (Compare with Fig. 3.)

3.2 LQR and state estimation

The standard LQ controller was augmented with integral states to provide servo following behaviour, [11]. A full six-state extended Kalman estimator was implemented despite the fact that two of the states, θ_1, θ_2, are measured directly. This provides some extra smoothing which is important given that the measured data is highly corrupted with noise. The matrix Riccati equations are solved iteratively, which while inefficient, has the advantage that the computations can be truncated if necessary but still deliver a near-optimal solution.

At higher elevation levels, the extreme nonlinearities in helicopter become evident in the widely varying elements of the $\mathbf{P}_\infty$ matrix and subsequent controller gains, $\mathbf{K}_\infty$ as shown in Fig. 5. It is interesting to note that an MPC controller was necessary to close the loop in Fig. 5 since using the wildly varying gain $\mathbf{K}_\infty$ lead to an overexcited unstable controller.

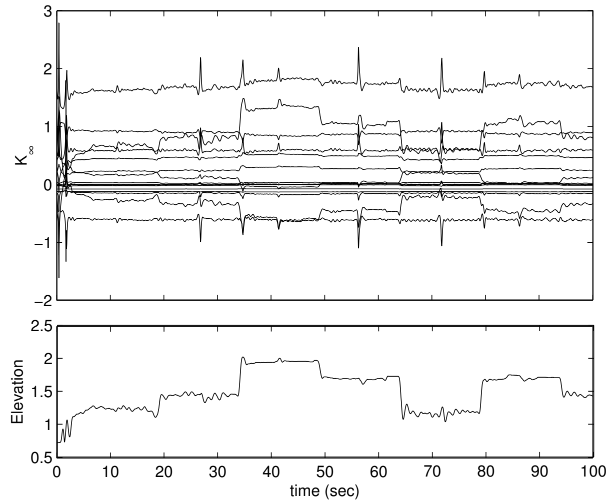

Figure 5: Variations in the 8 elements of the steady-state controller gain $\mathbf{K}_\infty$ due to plant nonlinearities and (lower) elevation trajectory.

If instead of computing the steady-state solution to the Riccati equation, $\mathbf{P}_\infty$, every time the model is updated (which in practice is every sample time), only one iteration of

$$\mathbf{P} \Leftarrow \mathbf{Q} + \mathbf{A}^T\mathbf{P}\left(\mathbf{R} + \mathbf{B}\mathbf{R}^{-1}\mathbf{B}^T\mathbf{P}\right)^{-1}\mathbf{A}$$

is performed. While suboptimal, this modification improves the robustness of the LQ controller. Still more suboptimal is to base the model re-linearisation on the setpoint as opposed to the current state. Since the setpoint changes far less frequently than the actual states, and that the dominant nonlinearity is the elevation, the result is that the gain follows that in Fig. 5 but without the destabilising spikes. This further improves the controlled performance.

3.2.1 LQG performance: Fig. 6 shows a typical servo response from the LQG controller updating $\mathbf{K}_\infty$ based on the model linearised about the setpoint. The controller could manage 600 iterations of the Riccati equation at a sample time of 50ms, although in practice never more than 200 were required.

One of the main drawbacks to optimal control design in state-space is the lack of guidelines that specifically address model robustness. It is clear that the poor performance at higher elevation levels in Fig. 6 is due somehow to a deficient model. This is evident from the trajectories of the adapting 'catch-all' gains that vary primarily as a function of elevation angle in the MPC implementation (not shown in Fig. 4). If the model was perfect, the gains would remain at 1. The MPC can cope with this by clever use of adaption, but our

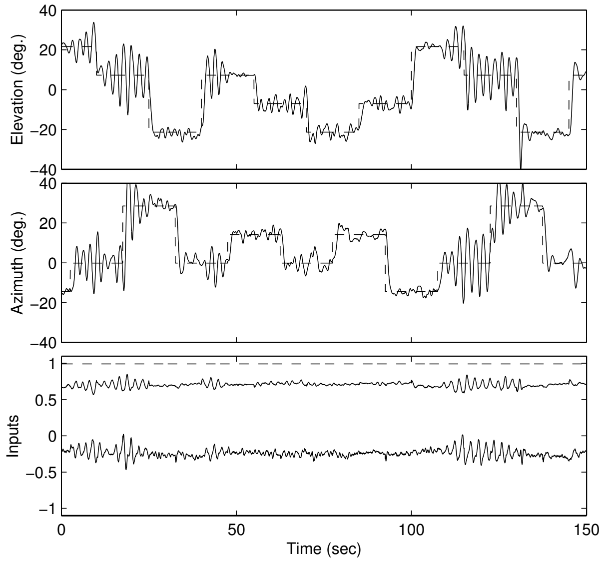

Figure 6: LQG controlled response. (Compare with Fig. 4.)

implementation of the optimal controller did not use it.

3.3 Optimal output control

Optimal output control (refer [3, Chapt 8] or what is sometimes known as parametric LQ control (PLQ), [12] is the control problem to find optimal constant feedback gains, $\mathbf{F}$, to minimise

$$\lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N-1} \left(\mathbf{x}(k)^T \mathbf{Q}\mathbf{x}(k) + \mathbf{u}(k)^T \mathbf{R}\mathbf{u}(k)\right) \tag{8}$$

subject to a linear stochastic plant model,

$$\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k + \mathbf{B}\mathbf{u}_k + \mathbf{w}_k, \quad \mathbf{y}_k = \mathbf{C}\mathbf{x}_k + \mathbf{v}_k \tag{9}$$

using *output* feedback,

$$\mathbf{u}_k = \mathbf{F}\mathbf{y}_k \tag{10}$$

Introducing the stationary state covariance matrix $\mathbf{P}$ as the solution to the discrete Lyapunov equation

$$\mathbf{P} = (\mathbf{A} + \mathbf{BFC})\, \mathbf{P}\, (\mathbf{A} + \mathbf{BFC})^T + \mathbf{R}_w + \mathbf{BFR}_v\mathbf{F}^T\mathbf{B}^T \tag{11}$$

and similarly the symmetric positive definite $\mathbf{S}$ as the solution to

$$\mathbf{S} = (\mathbf{A} + \mathbf{BFC})^T\, \mathbf{S}\, (\mathbf{A} + \mathbf{BFC}) + \mathbf{Q} + \mathbf{C}^T\mathbf{F}^T\mathbf{RFC} \tag{12}$$

then the optimum gain is given by

$$\mathbf{F} = -\left(\mathbf{B}^T\mathbf{SB} + \mathbf{R}\right)^{-1} \mathbf{B}^T\mathbf{SAPC}^T \left(\mathbf{CPC}^T + \mathbf{R}_v\right)^{-1} \tag{13}$$

where $\mathbf{R}_v, \mathbf{R}_w$ are the covariances of the measurement and state noise respectively. One way to compute $\mathbf{F}$ is to choose an initial stabilising gain $\mathbf{F}_0$ and then iterate around Equations 11–13 until convergence. However [12] caution that the convergence properties of this coupled equation system are not well understood and the solution is comparatively computationally expensive although our tests indicate that under 10 outer iterations usually suffice. Modifications to improve convergence have been proposed by various authors, see e.g. [4].

The structure of the PLQ with servo properties is given in Fig. 7 following that suggested in [3, §8.2]. To ease the online computation requirement, (which is an order of magnitude more than in the LQG case), the optimal gain $\mathbf{F}$ was pre-computed offline using the descent Anderson-Moore method, [12], in a 12-position lookup table using the elevation setpoint as the scheduling variable. The gain $\mathbf{K}_1$ was numerically optimised offline using the plant model from section 2.1 in the upper (unstable) elevation position.

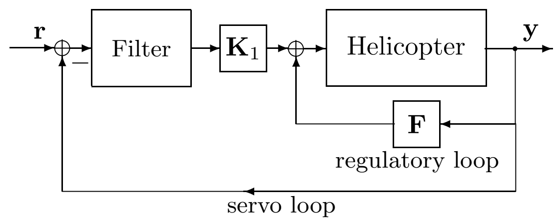

Figure 7: Structure of the PLQ controller

3.3.1 Output optimal control performance:

Fig. 8 shows that while the elevation response at angles below the horizontal is adequate, in the unstable region, it is barely stable despite the fact that the servo gain was tuned for this region. Disappointingly, the azimuth control is also poor, although a redesign of the servo filter perhaps using only PD rather than incorporating an integral term may solve this problem.

4 Performance comparison

A comparison of the controllers performance in terms of the integral of the absolute error (IAE) is given in Table 1 for both setpoint following, as shown in Figs 4 and 6, and disturbance rejection where the weight was pulsed by ±10% at times $t = 125$ and $t = 245$ as shown in Fig. 9. In the case of MPC, there is the possibility to use knowledge of future setpoint changes. This acausal modification has the potential to improve the performance of the servo response. The values presented in Table 1 are normalised by the acausal MPC case for the plant with the counter-weight in the nominal position.

Clearly from comparing Figs 4 and 6, MPC outperforms LQG. The variance is similar in both cases in

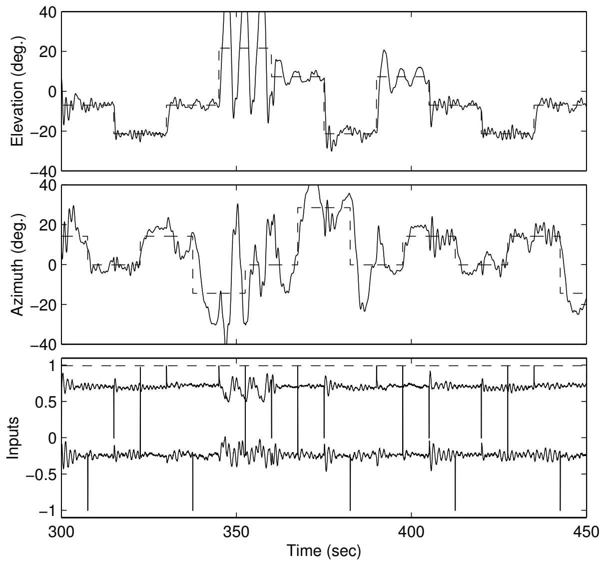

Figure 8: PLQ controlled response

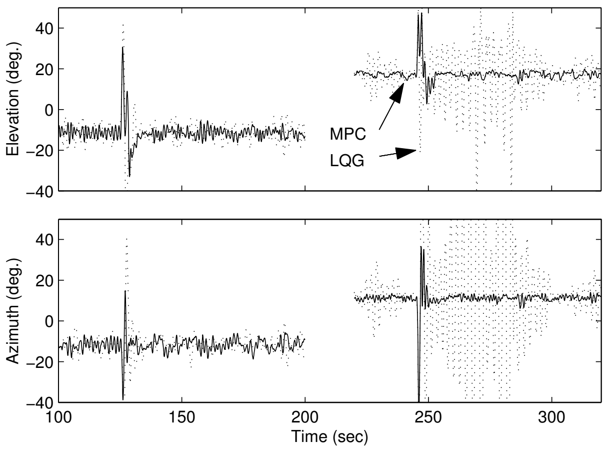

Figure 9: A comparison between MPC (solid) and LQG (dotted) for disturbance rejection at two distinct operating regions.

Fig. 9 up until the unmeasured counter-weight disturbance after which the MPC again outperforms the LQG. The LQG has problems in unstable upper elevation region. The MPC IAE performance criteria are around half that for LQG.

The PLQ response is generally poor, except in elevation below the horizontal. The IAE for disturbance rejection is less than in the LQG case indicating that the regulatory response due to the optimal output feedback is good, as opposed to the command-following behaviour. However this control methodology shows promise, and it is hoped with more careful tuning of the general filters one can improve the servo response.

5 Conclusions

Three very different optimal control methodologies were applied to a strongly-interacting, high-speed, nonlinear, unstable bench-scale helicopter. The application is non-trivial due to the need for frequent model re-linearisation, the demand for relatively short sampling times combined with substantial controller computation, and significant random disturbances.

The controlled response both to commands and to disturbances, introduced primarily by air turbulence, is superior to that using PID controllers. MPC outperformed the other, linear based optimal controllers, although it has substantial inter-sample computation demands and by far the largest memory requirements. LQG based on a model linearised every sample time proved satisfactory below the horizontal, but above, the instability of the open loop plant, combined with an increasing model/plant mismatch proved too much. One solution was to schedule the optimum gain based on the elevation angle. While sub-optimal, this proved sufficiently robust.

The optimal output based controller (PLQ) with servo-following capability avoided crashing, although it would be disingenuous to suggest that at present it be a serious candidate for any industrial application. However the response at low elevation angles is good, and there is promise that the poor azimuth performance could be addressed using a more careful controller design.

References

[1] David Wilson. Optimal control teaching using Matlab: Have we reached the turning point yet? In Tore Bjørnarå, editor, *Nordic Matlab Conference*, pages II–246 – II–251, Oslo, Norway, 17–18 October 2001. ISBN 82-995995-0-9.

[2] M. Morari and J.H. Lee. Model predictive control: Past, present and future. *Computers and Chemical Engineering*, 23:667–682, 1999.

[3] Frank L. Lewis and Vassilis L. Syrmos. *Optimal Control*. John Wiley & Sons, 2 edition, 1995.

[4] T. Rautert and E.W. Sachs. Computational design of optimal output feedback controllers. Technical Report Nr. 95-12, Universität Trier, FB IV, Germany, June 1995.

[5] Jonas Balderud and David Wilson. Application of Predictive Control to a Toy Helicopter. In *IEEE Conference on Control Applications*, Glasgow, Scotland, 18–20 September 2002. IEEE.

[6] Jonas Balderud. Modelling and Control of a Toy Helicopter. Master's thesis, Karlstad Univer-

Table 1: IAE performance comparisons for both servo and regulatory responses with the counter-weight in the nominal position and at 75% (deliberate model/plant mismatch). Values normalised by the nominal acausal MPC case.

Scheme	*Nominal*		*Model/plant mismatch*	
	Setpoint following	*Disturbance rejection*	*Setpoint following*	*Disturbance rejection*
PID	16.0	2.36		
MPC (acausal)	1	1	1.11	1.05
MPC (causal)	1.40	1.00	1.36	1.05
LQG	1.54	2.32	1.79	2.73
Output LQ	2.90	1.99	2.64	1.58

sity, Electrical Engineering Department, December 2001.

[7] David Clarke. *Advances in Model-Based Predictive Control*. Oxford University Press, 1994.

[8] P. Krauss, K. Dass, and H. Rake. Model-based predictive controller with Kalman filtering for state estimation. In D. Clarke, editor, *Advances in Model-Based Predictive Control*, pages 69–83. Oxford University Press, New York, 1994.

[9] C. R. Cutler and B.L. Ramaker. Dynamic Matrix Control — A Computer Control Algorithm. In *Proc. 1980 Joint automatic Control Conference*, San Francisco, 1980. American Institute of Chemical engineers.

[10] J.H. Lee and N.L. Ricker. Extended Kalman Filter Based Nonlinear Model Predictive Control. *Ind. Eng. Chem. Res.*, 33(6):1530–1541, 1994.

[11] Katsuhiko Ogata. *Discrete-Time Control Systems*. Prentice–Hall, 2 edition, 1995.

[12] Pertti M. Mäkilä and Hannu T. Toivonen. Computational Methods for Parametric LQ Problems – A Survey. *IEEE transactions on Automatic Control*, AC-32(8):658–671, 1987.