



This is a repository copy of *Towards Real-Time, Non-Linear Quadratic Optimal Regulation*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/82973/>

Monograph:

Harrison, R.F. and Banks, S.P. (1997) *Towards Real-Time, Non-Linear Quadratic Optimal Regulation*. Research Report. ACSE Research Report 693 . Department of Automatic Control and Systems Engineering

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

PS940047
X

Towards real-time, non-linear quadratic optimal regulation.

Robert F Harrison and Stephen P Banks

Department of Automatic Control and Systems Engineering,
The University of Sheffield,
Sheffield, UK.

Corresponding Author: Robert F Harrison,
Department of Automatic Control and Systems Engineering,
The University of Sheffield, Sheffield, S1 3JD, UK.
Tel. +44 (0)114 2225139
Fax. +44 (0)114 2780409
Email r.f.harrison@sheffield.ac.uk

Research Report No 693

Abstract

A novel non-linear design method based on linear quadratic optimal control theory is presented that applies to a wide class of non-linear systems. The method is easy to apply and results in a globally optimal solution that can be implemented in real-time. The key feature of the *design* method is the introduction of state-dependence in the weight matrices of the usual linear quadratic cost function, leading to a *non-linear* control policy, even for linear dynamics. To demonstrate the method, a simple vehicle suspension model with a cubic damping force is used, in conjunction with non-linear state penalties that better reflect the *engineering* objectives of active vehicle vibration suppression. A number of simulations is conducted and compared with a passively mounted vehicle.

Keywords: Non-linear; optimal control; matrix sign function; active suspension; parallel algorithms.

1. Introduction

Linear quadratic (LQ) optimal control theory is a highly developed approach for the synthesis of linear optimal control laws that has been widely applied. In particular, the infinite-time-horizon solution has appeal for the regulation of processes that are well modelled by linear time-invariant dynamics because the solution comprises a set of static gains that are calculated once, off-line, and are implemented causally thereafter. The finite-time-horizon solution, while being more general, and admitting of time-varying dynamics or weighting parameters, is essentially an off-line procedure because the associated (differential) Riccati equation must be solved backwards in time. Incorporation of such a scheme in a closed-loop, real-time process is not therefore possible. The receding horizon approach attempts to overcome this difficulty by repeatedly solving the open-loop, finite-time-horizon problem for short periods into the future, and using this as the feedback gain over the time step. It does this at the expense of optimality, which may or may not be important from the point of view of the practising engineer.

A further disadvantage of the LQ philosophy is that, being a linear feedback, the control signal is affected in the same way by small and large signals. In many applications it may be preferable to ignore small error signals (due, for instance, to measurement noise) as far as possible, while responding optimally to large errors. It may also be desirable to switch attention between control objectives depending on their current values. The receding horizon strategy

approximates the former behaviour to a certain extent, but is unable to address the latter.

For non-linear dynamics the situation is exacerbated, with few explicit solutions for their optimal quadratic control as yet known except those based on series expansions. These are unrealisable, unless via truncation, leading to a loss of optimality and possibly stability. Normally optimal quadratic control for non-linear systems is conducted numerically and tends, inherently, to be non-causal.

In this paper we make use of a new result that generalises the LQ theory to non-linear systems to provide a non-linear design method that overcomes some of the difficulties mentioned above. This non-linear quadratic (NLQ) method applies to systems having a broad class of non-linear dynamics with state-dependent weighting matrices (the design degrees of freedom). In brief, it turns out that the infinite-time-horizon LQ regulator problem, when solved afresh at every point on the state trajectory, leads to a globally optimal control policy [1]. For admissible system dynamics, the weighting parameters can be made to be functions of the state variables. Thus, in addition to handling non-linear dynamics, the design stage allows for the introduction of state-dependence in the weighting matrices, leading to a more flexible control strategy.

Our method is causal, but has considerable computational overhead. However, by using a solution to the Riccati equation based upon the matrix sign function [2], it is possible to derive a parallel algorithm [3] suitable for real-time implementation.

In order to demonstrate the approach we consider the active control of a vehicle suspension system in

200412445



order to reject disturbances induced by surface asperity: a problem that has received much attention in the past although usually for linearised dynamics. Clearly, the assumption of linearity may often be valid but some designs are inherently non-linear such as the oleopneumatic shock struts of aircraft landing gear.

While the LQ approach is attractive in that it is possible to penalise different variables so as to trade-off between, say, ride comfort and handling, or comfort and suspension travel, the way these variables are treated is essentially fixed – no provision is made to allow the suspension to distinguish between a smooth surface and a rough one. Evidently, while comfort might be a prime objective under normal circumstances, in extreme conditions the suspension should be stiffened to avoid hitting its limits, hence incurring damage. This is true even if the dynamics are linear up to this point. Although, in principle, time-varying weighting parameters are allowed in the (finite-time) LQ approach, lack of prior knowledge of the surface profile, and the acausal calculation for the solution makes the introduction of these difficult. The required amplitude dependence can never, therefore, be achieved through the LQ approach.

We illustrate our approach on a simple non-linear model incorporating a non-linear damping element, and compare our results with the passive system and the optimal NLQ solution using fixed weighting matrices.

The remainder of the paper is organised as follows. In §2, to motivate the work, the LQ regulator problem is first set out, the generalised results are stated and the robust solution of the infinite-time-horizon LQ problem is outlined. In §3 the suspension model is presented and the choice of design parameters is discussed. The results of a series of experiments are described in §4, followed by conclusions in §5.

2. The design method

2.1. Linear quadratic regulator

The LQ optimal regulation problem is expressed as follows:

minimise the cost function:

$$J = \int_0^{\infty} (\mathbf{x}' Q \mathbf{x} + \mathbf{u}' R \mathbf{u}) dt \quad (1)$$

subject to the linear time invariant dynamics:

$$\dot{\mathbf{x}} = A \mathbf{x} + B \mathbf{u} \quad (2)$$

where \mathbf{x} is an n -vector of system states, \mathbf{u} is an m -vector of control variables, A and B are matrices of appropriate dimension and the superscript, t , indicates transposition. The matrices Q and R are

positive semi-definite and definite, respectively, and are used to penalise particular states and controls according to the engineering objective.

It is well known (e.g. [8]) that the control policy which solves the above optimisation problem is a linear combination of the system states and is given by:

$$\mathbf{u} = K \mathbf{x} \quad (3)$$

where K is in turn given by:

$$K = -R^{-1} B' P \quad (4)$$

and P is the positive definite solution of the algebraic matrix Riccati equation:

$$0 = PA + A' P - P B R^{-1} B' P + Q \quad (5)$$

A unique, positive definite solution to the above exists if the pair (A, B) is stabilizable and (A, Γ) is detectable, with $Q = \Gamma' \Gamma$.

2.2. Global, non-linear quadratic regulator

The extension of the above to non-linear systems looks identical, except that, instead of performing a single optimisation and applying the resulting gain-matrix for all time, the optimisation has to be carried out at every time step. Consider a non-linear dynamical system that can be expressed in the form:

$$\dot{\mathbf{x}} = A(\mathbf{x}) \mathbf{x} + B(\mathbf{x}) \mathbf{u} \quad (6)$$

where the Jacobians of $A(\mathbf{x})$ and $B(\mathbf{x})$ are subject to some bounded growth conditions (Lipschitz), then at each point, $\bar{\mathbf{x}}$, on the state trajectory, a linear system is defined with fixed $A = A(\bar{\mathbf{x}})$ and $B = B(\bar{\mathbf{x}})$. In [1]

it is shown that solving the infinite-time LQ optimal control problem, point-wise on the state trajectory, results in the globally optimal, stabilising quadratic control policy for systems described by equation (6). Thus, by choosing the \mathbf{u} that minimises the usual quadratic cost function at every time step, we have a globally optimal control policy for a very wide class of non-linear systems. Evidently, $A(\bar{\mathbf{x}})$, $B(\bar{\mathbf{x}})$ and Q are subject, point-wise, to the same conditions as for the linear case. It is clear that the proposed solution is identical to the one obtained from equations (3, 4 and 5) when the dynamics are linear.

As an aside, the dual situation follows directly from the reasoning in [1] and thus state estimation is possible via a non-linear observer although this aspect is not dwelt on here.

Because the control synthesis takes place point-wise, the designer is now free to select Q and R in ways which are more directly applicable to the control engineering objectives. In particular, these can be made to be functions of the instantaneous state variables, i.e.

$$\begin{aligned}\ddot{y}_1 &= -\frac{k_1}{m_1}(y_1 - y_2) - \frac{c_1}{m_1}(\dot{y}_1 - \dot{y}_2) \\ &\quad - \frac{c_2}{m_1}(\dot{y}_1 - \dot{y}_2)^3 + \frac{1}{m_1}f \\ \ddot{y}_2 &= \frac{k_1}{m_2}(y_1 - y_2) + \frac{c_1}{m_2}(\dot{y}_1 - \dot{y}_2) \\ &\quad + \frac{c_2}{m_2}(\dot{y}_1 - \dot{y}_2)^3 - \frac{k_2}{m_2}(y_2 - d) - \frac{1}{m_2}f\end{aligned}\quad (10)$$

We choose state variables thus:

$x_1 = y_1, x_2 = \dot{y}_1, x_3 = y_2, x_4 = \dot{y}_2$, and identify the control signal, u , with the force, f . Evidently equations (10) can be put into the required form (6), and the Jacobians of both A and B are clearly Lipschitz.

3.1. Design objectives

For the purposes of this paper let us suppose that our primary objective is to minimise passenger discomfort. We do this by attempting to reduce the accelerations to which the passenger is subject – vertical only, in this simple case. Thus a candidate for the cost function is $\ddot{y}_i = C_a(\bar{x})\mathbf{x} + D_a(\bar{x})\mathbf{u}$, where $C_a(\bar{x})$ is the second row of $A(\bar{x})$ and $D_a(\bar{x})$ is the second element of $B(\bar{x})$. However, ride comfort can only take precedence when safety and integrity are not compromised. Thus it is necessary to penalise some measure which embodies these ideas, usually via the ‘‘rattlespace deflection’’,

$y_1 - y_2 = C_r\mathbf{x} = [1 \ 0 \ -1 \ 0]\mathbf{x}$. In the conventional

LQ approach we construct a cost function thus:

$$\begin{aligned}J &= \int_0^{\infty} (q_a \ddot{y}_i^2 + q_r (y_1 - y_2)^2 + ru^2) dt \\ &= \int_0^{\infty} \left(\mathbf{x}' (q_a C_a' C_a + q_r C_r' C_r) \mathbf{x} \right. \\ &\quad \left. + 2\mathbf{x}' q_a C_a' D_a u + (q_a D_a^2 + r) u^2 \right) dt\end{aligned}\quad (11)$$

Letting $N = q_a C_a' D_a$ and $R = q_a D_a^2 + r$ we accommodate the cross-term in the usual way [8] thus:

$Q \leftarrow Q - NR^{-1}N'$, $A \leftarrow A - BR^{-1}N'$ with the original $Q = q_a C_a' C_a + q_r C_r' C_r$. q_a, q_r are used to control the trade-off between ride and handling. The new Q and A are now used in the standard equations (6) and (7).

In the non-linear design procedure we make q_r state-dependent, thus: $q_r(\bar{x}) = 500\psi(y_1 - y_2, .02, .001)$ with

$$\psi(\xi, \theta, \delta) = \begin{cases} ((\xi - \theta) / \delta)^4, & \xi > \theta \\ 0, & |\xi| \leq \theta \\ ((\xi + \theta) / \delta)^4, & \xi < -\theta \end{cases}\quad (12)$$

where $\theta \geq 0$ defines a dead-zone and $\delta > 0$, the distance within which ψ first reaches unity. The rationale for this functional form is as follows. The primary objective is to reduce body acceleration hence the constant $q_a = 10000$. The secondary objective, which can over-ride the first for safety reasons, is to reduce overly large excursions in the suspension strut. Thus, for a rattlespace of $\pm 0.055\text{m}$, a dead-zone of $\pm 0.02\text{m}$ is allowed before control action is taken; the non-linearity increasing to unity within the next 0.001m of travel and dominating the cost function very rapidly as the limits are approached. We have been guided here by the function chosen in [7], however, any other suitable function is a candidate. In addition, acceleration or other variables could equally well be weighted in this way. For comparison, we also consider a conventionally weighted rattlespace deflection with $q_r = 1000$. $r = 0.0001$ throughout.

We use the passive system (equation (10) with $f(t) = 0$ for all t) as a reference and compare its behaviour with that of the actively controlled models for a surface profile [7] given by

$$d(t) = \begin{cases} a(1 - \cos 8\pi t), & 0 \leq t \leq 0.25 \\ 0, & \text{otherwise} \end{cases}\quad (13)$$

where a is one-half the height of the hump.

4. Results

Each of figures 2 onwards shows the passive, conventionally weighted (CW) and state-dependently weighted (SDW) behaviour of the vehicle model. These are indicated by a dashed, a dotted and a solid curve, respectively. Figures 2–5 relate to a modest disturbance of maximum height 2cm which does not threaten an approach towards the rattlespace limits. Here, because $|\dot{y}_1 - \dot{y}_2| < 0.02$ for all time, the rattlespace deflection is never penalised, and because acceleration dominates strongly in the cost function for both CW and SDW, it is not possible to distinguish between these behaviours. Figures 6–9 show the behaviour for a much more severe disturbance, with a maximum height of 11cm that forces the rattlespace weighting into play.

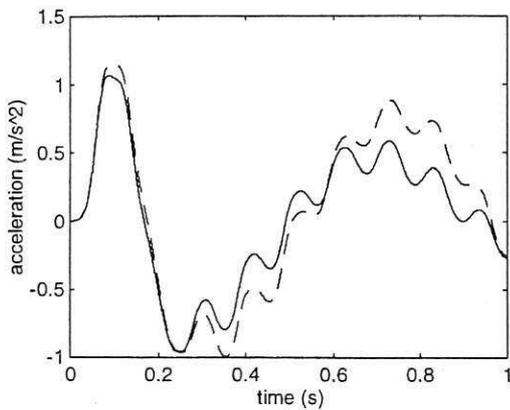


Figure 2 Acceleration vs. time, $a=0.01$.

Figures 2, 3 and 4 indicate the improvements attained by both CW and SDW, while figure 5 hints that these come at the expense of marginally increased unsprung mass displacement.

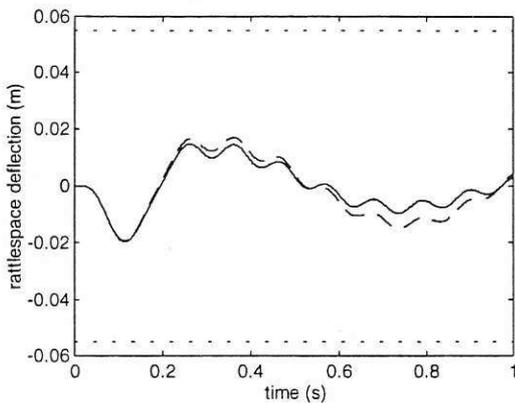


Figure 3 Rattlespace deflection vs. time, $a=0.01$.

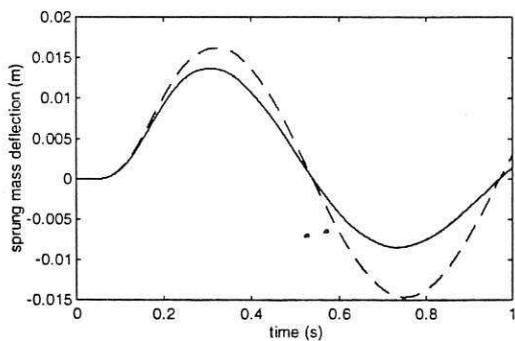


Figure 4 Sprung mass deflection vs. time, $a=0.01$.

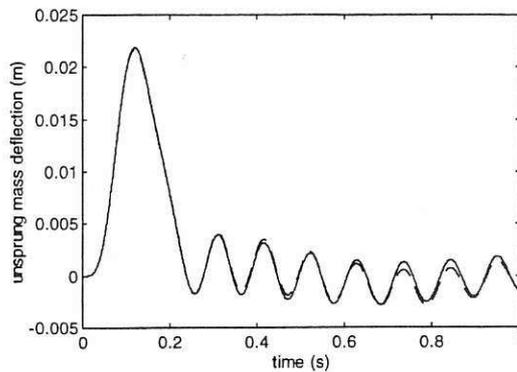


Figure 5 Unsprung mass deflection vs. time, $a=0.01$.

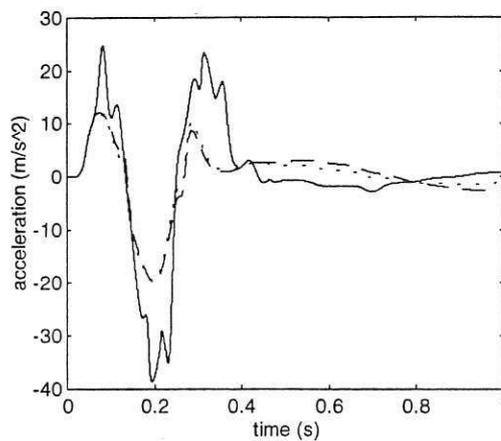


Figure 6 Acceleration vs. time, $a=0.55$.

Figure 6 shows how, for a severe disturbance the transient acceleration is degraded while tending towards an improved steady-state. Although the acceleration is penalised conventionally, there is an effect when the SDW for rattlespace becomes operative ($\sim 0.06s$). This sacrifices comfort for safety leading to a response which is worse than the passive system for the first $\sim 0.5s$. There is little difference between the passive and CW systems.

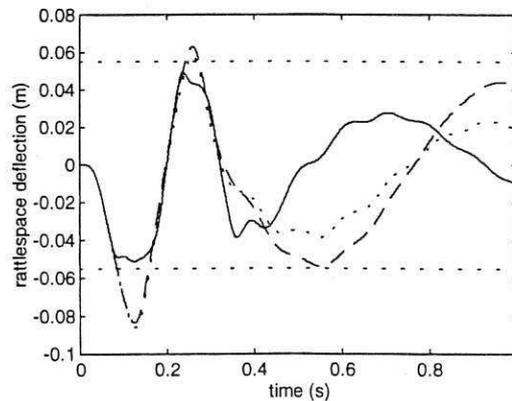


Figure 7 Rattlespace deflection vs. time, $a=0.55$.

Figure 7 highlights a more important difference in rattlespace behaviour. Again CW and SDW deliver

improvements over the passive behaviour but now the value of SDW is evident. The CW system exceeds the limits of travel in the suspension strut – “bottoming” occurs – while SDW rapidly counteracts the approach (albeit for a deterioration in ride). Nonetheless, only by de-tuning the CW system can bottoming be avoided, and this could not be guaranteed.

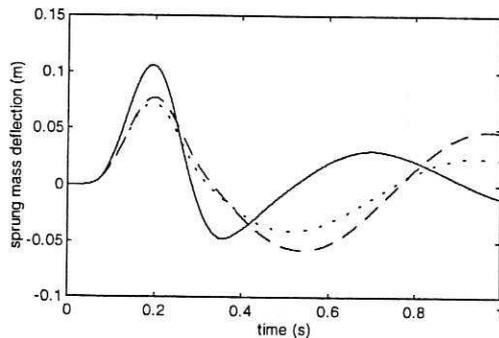


Figure 8 Sprung mass deflection vs. time, $a=0.55$.

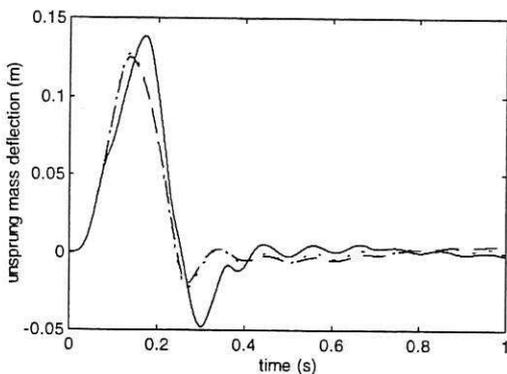


Figure 9 Unsprung mass deflection vs. time, $a=0.55$.

Figures 8 and 9 indicate that improvements in ride and handling may come at the expense of increased excursions of the sprung and unsprung masses themselves. However, in a fuller design it would be straightforward to penalise these too.

5. Conclusions

A new method for the design and synthesis of globally optimal, non-linear control laws is proposed, based on a generalisation of LQ optimal control theory. The method is simple to apply and affords greater design flexibility (SDW) than the conventional approach (CW). The resulting controller is non-linear, even for linear dynamics, and can be implemented in real-time. To illustrate the method a simple two-degree-of-freedom quarter-car model with cubic damping has been studied using a non-linear penalty function. Preliminary results show that the method has applicability and could easily be tuned to provide desirable closed-loop behaviour. Some areas for future work are the

selection of more appropriate penalty functions, $Q(x)$, $R(x)$, and the real-time implementation of such a system.

NB In the final paper a fuller analysis of the results would be presented.

References

- [1] Banks, S.P. and Mhanna, K.J., Optimal control and stabilisation for non-linear systems, *IMA Journal of Mathematical Control and Information* (1992) **9**, 179–196.
- [2] Gardiner, J.D. and Laub, A.J., A generalisation of the matrix-sign-function for algebraic Riccati equations, *International Journal of Control* (1986) **44**, 823–832.
- [3] Gardiner, J.D. and Laub, A.J., Parallel algorithms for algebraic Riccati equations, *International Journal of Control* (1991) **54**, 1317–1333.
- [4] Laub, A.J., A Schur method for solving algebraic Riccati equations, *IEEE Transactions on Automatic Control* (1979) **24**, 913–921.
- [5] Roberts, J.D., Linear model reduction and solution of the algebraic Riccati equation by use of the sign function, *International Journal of Control* (1980) **32**, 677–687.
- [6] Bunch, J.R., Analysis of the diagonal pivoting method, *SIAM Journal on Numerical Analysis* (1971) **8**, 656–680.
- [7] Lin, J-S. and Kanellakopoulos, I., Non-linear design of active suspensions, *IEEE Control Systems* (1997) **17**, 45–59.
- [8] Friedland, B., *Control system design: an introduction to state-space methods*, McGraw-Hill Book Co., New York, 1987.

