ELSEVIER

Systems & Control Letters 26 (1995) 223–230

SYSTEMS & CONTROL LETTERS

Finite- dimensional optimal controllers for nonlinear plants

John B. Moore [a,*,1], John S. Baras [b]

[a] Department of Systems Engineering, Research School of Information Sciences and Engineering, Australian National University, Canberra, ACT 0200, Australia

[b] Department of Electrical Engineering, University of Maryland, College Park, MD 20742, USA

Received 27 November 1994; revised 12 May 1995

Abstract

Optimal risk sensitive feedback controllers are now available for very general stochastic nonlinear plants and performance indices. They consist of nonlinear static feedback of so called information states from an information state filter. In general, these filters are linear, but infinite dimensional, and the information state feedback gains are derived from (doubly) infinite dimensional dynamic programming. The challenge is to achieve optimal finite dimensional controllers using finite dimensional calculations for practical implementation.

This paper derives risk sensitive optimality results for finite-dimensional controllers. The controllers can be conveniently derived for 'linearized' (approximate) models (applied to nonlinear stochastic systems). Performance indices for which the controllers are optimal for the nonlinear plants are revealed. That is, inverse risk-sensitive optimal control results for nonlinear stochastic systems with finite dimensional linear controllers are generated. It is instructive to see from these results that as the nonlinear plants approach linearity, the risk sensitive finite dimensional controllers designed using linearized plant models and risk sensitive indices with quadratic cost kernels, are optimal for a risk sensitive cost index which approaches one with a quadratic cost kernel. Also even far from plant linearity, as the linearized model noise variance becomes suitably large, the index optimized is dominated by terms which can have an interesting and practical interpretation.

Limiting versions of the results as the noise variances approach zero apply in a purely deterministic nonlinear H_∞ setting. Risk neutral and continuous-time results are summarized.

More general indices than risk sensitive indices are introduced with the view to giving useful inverse optimal control results in non-Gaussian noise environments.

Keywords: Nonlinear control; Optimal controllers; Risk-sensitive control; Information state filters; Finite-dimensional controllers

1. Introduction

Optimal open loop control of nonlinear plants goes back many decades, and for this length of time there has been the challenge to develop a feedback optimal control theory so as to exploit the robustness of feedback for nonlinear control in an optimal setting. Such results have been achieved for very general nonlinear plants and indices in a stochastic setting, see [10]. The catch is that the controllers are infinite dimensional and require infinite dimensional calculations for both design and implementation. Also, not many engineers are comfortable working in a stochastic setting for nonlinear signal models. The controllers actually consist of a linear infinite-dimensional information state filter and static nonlinear feedback of the information states. The latter is derived by an off-line (doubly) infinite dimensional dynamic programming equation.

* Corresponding author. E-mail: john.moore@anu.edu.au.

[1] The author acknowledges the support of the Cooperative Research Centre for Robust and Adaptive Systems at the ANU.

SSDI 0167-6911(95)00025-9

A follow on result in [12] develops risk sensitive generalizations of the earlier risk neutral theory so as to tune robustness with a so-called risk sensitive parameter, and moreover in this setting, achieves limiting results as the noise variances become zero which give optimal nonlinear feedback controllers for deterministic nonlinear plants. The small noise limit optimizing algorithms are differential games which allow tuning for robustness, and indeed can be interpreted as optimal in a worst case deterministic noise setting. They yield a natural nonlinear generalization of the so-called H_∞ controllers. Follow-on direct deterministic derivations of the latter are given in [11]. Earlier direct directions of H_∞ controllers are given in [2, 5, 14, 16–18]. Again however, the controllers are infinite dimensional in general. For background reading see books [9, 19].

In order to achieve finite-dimensional optimal controllers exploiting the power of the recent theory, first efforts imposed certain restrictions on the plant. In [1, 7], the plant is assumed to be linear in the states, but not the controls, and the index kernel is assumed to be quadratic in the states but not in the controls. This results in a finite dimensional information state filter, but the static information state feedback control law is calculated off line by (singly not doubly) infinite dimensional dynamic programming equations. Of course, if the plant is additionally linear in the controls, and the index kernel is additionally quadratic in the controls, then this theory leads to known linear quadratic (risk sensitive) controllers [8, 20]. These controllers are of course linear and have the same dimension as the plant.

More recently, finite dimensional information filters for the continuous-time, risk sensitive, stochastic control setting have been achieved for a special class of nonlinear plants and indices, see [6]. Small noise limits for this setting are studied in [7]. The key idea is that the special nonlinearities in the plant must be absorbed in some way in the cost index. It is our observation that the absorption is not as general as it might be, perhaps because the authors limited themselves by the objective of generalizing the Beneš continuous-time filter results to this setting. Actually, the Beneš filter is equivalent, to within a coordinate basis, the Kalman filter as pointed out in [3].

In this paper, motivated by computer implementation requirements, we first generate discrete-time, optimal risk sensitive control results achieving (linear) finite dimensional information filters and (nonlinear static) information state feedback via (singly) infinite dimensional dynamic programming). The theorems do not restrict the plant nonlinearity but rather include terms in the performance index to compensate for these nonlinearities. The controller turns out to be identical to that derived for a model linearized with respect to the states, and a risk sensitive index with a kernel quadratic in the states. We go on to study even more restricted indices to achieve linear information state feedback and consequently linear, finite dimensional, controllers. Here the controller is identical to that derived for a model linearized with respect to both the states and the control inputs, and a risk sensitive index with a kernel quadratic in both the states and the controls. Since we give indices for which practical controllers are optimal for nonlinear plants, the results constitute so-called inverse (risk sensitive) optimal control theory.

Another motivation in our work is to demonstrate rigorously that linear finite-dimensional controllers can be optimal for nonlinear, but 'nearly' linear plants and a reasonable class of risk sensitive indices. Also, for nonlinear plants that can only be approximated reasonably by models linear in the states, appropriate optimal controllers can be designed which have finite dimensional linear dynamics, but requiring static nonlinear information state feedback designed from off-line infinite-dimensional dynamic programming calculations (or finite dimensional approximations to these).

Two sets of additional results are derived of more general interest. The first set concerns nonlinear plants driven by non-Gaussian noise. We show that finite dimensional controllers, as discussed above, can be achieved by working with more general risk sensitive indices than those involving simply the product of exponentials of costs at each time instant. Other functions can replace the exponentials to advantage.

The second set of additional results concerns the common practice in design to achieve robustness to unmodelled dynamics. One assumes more noise variance and initial state uncertainty variance than is really there in a plant. In particular, we show in the most general nonlinear design case, what is well known in the linear quadratic Gaussian case, namely that the cost index actually optimized in the design is a reasonable upper bound for the cost index in the absence of uncertainty.

Continuous-time versions of some of our results are stated without derivations, since these parallel closely the discrete-time case. Likewise risk neutral, and deterministic small noise limit versions of some of the results are only summarized without fleshing out all details, since these do not require additional insights.

In Section 2 we set up the risk sensitive control problem to lead into the new results for finite dimensional optimal controllers in Section 3. Conclusions are drawn in Section 4.

2. The risk sensitive control problem

2.1. The plant model

Consider the state space stochastic nonlinear model

$$\begin{aligned} x_{k+1} &= A(u_k)x_k + a(u_k,x_k) + B(u_k)u_k + w_k, \\ y_{k+1} &= Cx_k + c(x_k) + v_k, \end{aligned} \tag{1}$$

on $k = 0,1,2,\ldots,M$ with states x_k an n-vector and y_k an m-vector. The noise terms w and v are assumed i.i.d. with distributions $\phi(\cdot)$ and $\psi(\cdot)$, respectively, and perhaps control and state dependent. The matrix functions $A(\cdot), B(\cdot), C(\cdot)$ may be time varying but here we suppress a k subscript for notational simplicity.

A linearized approximation to the above model is assumed to be (1) with the nonlinear terms $a(\cdot,\cdot)$ and $c(\cdot)$ set to zero, namely the system

$$\begin{aligned} x_{k+1} &= A(u_k)x_k + B(u_k)u_k + \bar{w}_k, \\ y_{k+1} &= C(u_k)x_k + \bar{v}_k, \end{aligned} \tag{2}$$

with initial state $\bar{x}_0$. Here $\bar{w}$ and $\bar{v}$ are not necessarily assumed to be identical to w, v, but are assumed to be i.i.d. with densities $\bar{\phi}$ and $\bar{\psi}$, respectively. Typically, the covariances of the densities are chosen to be 'greater' than for the original model to compensate for neglecting the nonlinearites $a(.,.), c(.)$ in the linearization process. Of course, the linearization giving rise to (2) is only with respect to the states x_k, and a further linearization with respect to the controls leads to a linear model, denoted as

$$\begin{aligned} x_{k+1} &= A_k x_k + B_k u_k + \bar{w}_k, \\ y_{k+1} &= C_k x_k + \bar{v}_k. \end{aligned} \tag{3}$$

For the first new results of the paper, so as to keep fully within the standard risk sensitive optimal control framework, we require Gaussian noise and initial state uncertainty, as

$$\begin{aligned} & w_k \sim \mathcal{N}[0, Q_k(x_k,u_k)], \\ & v_k \sim \mathcal{N}[0, R(x_k,u_k)], \quad x_0 \sim \mathcal{N}[0,P_0], \\ & \bar{w}_k \sim \mathcal{N}[0,\bar{Q}_k], \\ & \bar{v}_k \sim \mathcal{N}[0,\bar{R}_k], \quad \bar{x_0} \sim \mathcal{N}[0,\bar{P}_0]. \end{aligned} \tag{4}$$

As stated earlier, a typical selection would be with

$$\bar{Q}_k \geqslant Q_k, \qquad \bar{R}_k \geqslant R_k, \quad \bar{P}_0 \geqslant P_0 \tag{5}$$

for all k. Also, for the limiting results as the noise variances approach zero, we require the following conditions on the nonlinearities.

The nonlinearities $a(\cdot,\cdot)$ and $c(\cdot)$ are uniformly continuous in x and bounded by an affine function of the norm of x. (6)

2.2. Cost functions

The risk sensitive cost function for admissible controls in the set $U_{0,M}$ is

$$J(u) = E\left[\exp\frac{\mu}{\varepsilon}\left(\sum_{i=0}^{M-1} L(x_i,u_i) + \Phi(x_M)\right)\right], \tag{7}$$

where $\theta = \mu/\varepsilon$ is a risk sensitive adjustment parameter. Here ε is a parameter which goes to zero in a small noise limit case, discussed later. For each ε chosen, μ must be selected suitably small for the index and associated information state filters to exist. In the limit as μ goes to zero (for fixed ε), then the exponential operation becomes linear and, the so-called, risk neutral control is achieved.

The control problem is to find u^* in $U_{0,M}$ such that

$$J(u^*) = \inf_{u\in U_{0,M}} J(u). \tag{8}$$

Here L is scalar and is nonnegative, and for the small noise limit results of this paper, is constrained as

The kernel $L(\cdot,\cdot)$ is uniformly continuous and bounded by a quadratic in the norms of x and u. (9)

Our aim is to restrict the class of $L(\cdot,\cdot)$ so as to achieve finite dimensional controllers, and then decide when the restricted class of indices has practical value.

Other cost function formulations than (7) turn out to be appropriate to consider in the case of non-Gaussian noise models, as in (1) when $\phi(\cdot)$ and $\psi(\cdot)$ are non-Gaussian. In this case, indices which incorporate the function $\phi(\cdot)$ and $\psi(\cdot)$ can be useful. Thus we introduce the index

$$J(u) = E\left[\prod_{i=0}^{M-1} f(L(x_i,u_i)) f(\Phi(x_M))\right], \tag{10}$$

where $f(\cdot)$ is a scalar function of its argument. Clearly, (10) collapses to (7) when $f(\cdot)$ is $\exp(\cdot)$.

2.3. A measure change

For the model (1) defined on $(\Omega, \mathscr{F}, P)$ with $\mathscr{G}_k = \sigma(x_l, y_l;\ 0 \leqslant l \leqslant k)$ and $\mathscr{Y}_k = \sigma(y_l,\ 0 \leqslant l \leqslant k)$, consider the measure change from P to $\bar{P}$ as

$$\Lambda_{0,k} := \left.\frac{\delta P}{\delta \bar{P}}\right|_{\mathscr{G}_k} = \prod_{i=0}^{k} \frac{\phi(y_{i+1} - Cx_i - c(x_i))}{\phi(y_{i+1})}. \tag{11}$$

This yields y_k i.i.d under $\bar{P}$ as in [9].

2.4. Information state and adjoint process

Consider an information state associated with the model (1) and measure change (11) as $q_{k|k}(x)$ satisfying the defining equation for all $b : \mathbb{R}^n \to \mathbb{R}$ Borel test functions

$$\left\langle b, \exp\left\{\frac{\mu}{\varepsilon}\left(\sum_{i=0}^{k-1} L(x_i, u_i)\right)\right\} q_{k|k}\right\rangle$$
$$:= \int_{\mathbb{R}^n} b(x) \exp\left\{\frac{\mu}{\varepsilon}\left(\sum_{i=0}^{k-1} L(x_i, u_i)\right)\right\} q_{k|k}(x)\mathrm{d}x$$
$$= \bar{E}\left[\Lambda_{0,k} \exp\left\{\frac{\mu}{\varepsilon}\left(\sum_{i=0}^{k-1} L(x_i, u_i)\right)\right\} b(x_k) \mid \mathscr{Y}_k\right], \tag{12}$$

where $\bar{E}$ is the expectation under $\bar{P}$. As shown in [12], the information state satisfies the linear integral equation

$$q_{k+1|k+1}(x) = \frac{1}{\phi(y_{k+1})} \int_{\mathbb{R}^n} \phi(y_{k+1} - C\xi - c(\xi)) \times \psi(x - A(u_k)\xi - B(u_k)u_k - a(\xi, u_k)) \times \exp\left(\frac{\mu}{\varepsilon} L(\xi, u_k)\right) q_{k|k}(\xi)\mathrm{d}\xi. \tag{13}$$

In shorthand (linear operator) notation

$$q_{k+1|k+1} = \Sigma(u_k, y_{k+1}) q_{k|k}. \tag{14}$$

An adjoint process is defined from a backwards recursion

$$\beta_{k-1|k-1} = \Sigma(u_{k-1}, y_k) \beta_{k|k}. \tag{15}$$

2.5. Dynamic programming

As shown in [12], minimization of the risk sensitive control index is equivalent to the following minimization

$$S(q,k) = \inf_{u \in U_{k,M-1}} \bar{E}[\langle q_{k|k}, \beta_{k|k}\rangle \mid q_{k|k} = q]. \tag{16}$$

Moreover, $S(\cdot,\cdot)$ satisfies the dynamic programming equation

$$S(q,k) = \inf_{u \in U_{k,k}} \bar{E}[S(\Sigma(u, y_{k+1})q, k+1)]$$
$$S(q,M) = \left\langle q, \exp\left(\frac{\mu}{\varepsilon}\Phi\right)\right\rangle. \tag{17}$$

A key observation for one of the results to follow is that the above results of [12] can be generalized for coping with the more general cost indices (10) simply by replacing $\exp(\cdot)$ in the above equations (12)–(17) by $f(\cdot)$.

3. Finite dimensional optimal controllers

3.1. Exploiting linearization with respect to states

From (12), we see that under the Gaussian assumptions (4), then $q_{0|0} \sim N[0, P_0]$ is Gaussian. Moreover, $q_{k|k}$ is Gaussian for all k if and only if the integrand of (13) is Gaussian with exponent quadratic in ξ, x. That is,

$$L(x_k, u_k) = (x_k' M(u_k) x_k + N(u_k) x_k + m(u_k)) \tag{18}$$

for some possibly time-varying $M(\cdot), N(\cdot)$ and $m(\cdot)$ of appropriate dimension. It follows that in this case, the information filter and adjoint filter are not only linear as in the general case, but finite dimensional of dimension n, in general.

The first key contribution of this paper is to observe that with appropriate selections of the cost index kernal $L(\xi, \cdot)$, then the plant nonlinearities $c(\xi)$ and $a(\xi, \cdot)$ can be 'absorbed' by the $L(\xi, \cdot)$ to achieve a finite dimensional (linear) information state filter, and information state feedback controller as in the following theorem.

Theorem 3.1. *Consider the discrete-time stochastic nonlinear control system* (1) *and risk sensitive cost index*

$$J(u) = E\left[\exp\left\{\frac{\mu}{\varepsilon}\left(\sum_{i=0}^{M-1} \bar{L}(x_{i+1}, x_i, u_i) + \Phi(x_M)\right)\right\}\right] \tag{19}$$

where the cost index kernel is mildly generalized to include x_{k+1} dependence and restricted as

$$\bar{L}(x_{k+1},x_k,u_k)$$

$$= \frac{\varepsilon}{\mu}\Big\{\|y_{k+1} - Cx_k - c(x_k)\|^2_{\bar{R}}$$

$$+\|x_{k+1} - A(u_k)x_k - B(u_k)u_k - a(x_k,u_k)\|^2_{\bar{Q}}$$

$$-\|y_{k+1} - Cx_k\|^2_{\bar{R}} - \|x_{k+1} - A(u_k)x_k - B(u_k)u_k\|^2_{\bar{Q}}\Big\}$$

$$+L(x_k,u_k) \tag{20}$$

for $L(x_k,u_k)$ given in (18), *where $\|l\|_X := l'Xl$. Then the linear information state filter* (13) *and information state feedback controller given from* (16)–(17) *has finite dimensional dynamics. Moreover, they are identical (within a scaling factor) to that for the plant linearized with respect to x, as in* (2), *and with the cost term* (7) *under* (18).

Proof. That the information states of the two control problems are identical (within a scaling factor) follows since (13) under (20) is

$$q_{k+1|k+1}(x) = \frac{1}{\mathcal{N}(y_{k+1},\bar{R})}\int_{\mathbb{R}^n}\mathcal{N}\left(y_{k+1} - c(u_k)\xi,\bar{R}\right)$$

$$\times\mathcal{N}\left(x - A(u_k)\xi - B(u_k)u_k,\bar{Q}\right)$$

$$\times\exp\Big\{\frac{\mu}{\varepsilon}(x_k'M(u)x_k + N(u_k)x_k$$

$$+ m(u_k))\Big\}q_{k|k}(\xi)\,\mathrm{d}\xi.$$

That this information state is given by finite dimensional dynamics follows since the exponents in the formulation are quadratic in x, ξ, see also [1]. □

3.2. Robustness properties

The above theorem is an inverse optimal control result in that it gives an index for which a class of (desirable) finite-dimensional controllers is optimum when applied to a nonlinear stochastic plant. In order to proceed with generalizations, specializations, and follow-on results, we work with a simplified version of (20), namely

$$\bar{L}(\cdot) = \frac{\varepsilon}{\mu}\Big\{(\|v_k\|^2_{\bar{R}} + \|w_k\|^2_{\bar{Q}}) - (\|v_k + c(x_k)\|^2_{\bar{R}}$$

$$+\|w_k + a(x_k,u_k)\|^2_{\bar{Q}})\Big\} + L(x_k,u_k). \tag{21}$$

From this equation we see that for near linear plants, in which the nonlinearity norms $\|a(\cdot,\cdot)\|, \|c(\cdot)\|$ are relatively insignificant, then

$$\bar{L}(\cdot) \simeq L(\cdot). \tag{22}$$

Also observe that under (5) and with the nonlinearities $a(\cdot,\cdot), c(\cdot) = 0$, then

$$\bar{L}(\cdot) \geqslant L(\cdot), \qquad \bar{J}(\cdot) \geqslant J(\cdot). \tag{23}$$

Consequently, consider the case of plants linearized with respect to the states and driven by Gaussian noise, and optimally controlled according to a risk sensitive index with nonnegative kernel quadratic in the states, but assuming more noise variance and initial state uncertainty than is there. Then the finite dimensional controller so derived is indeed optimal for an index which is an upper bound on the design index. This result also holds for more general indices and plants as follows. Consider the nonlinear stochastic plant (1) driven by Gaussian noise, but where for design purposes, the plant noise variances and initial state uncertainty is assumed greater than is actually the case as in (5). Consider also that the optimal controller for the index (7) is applied (being infinite dimensional in general). Then a straightforward generalization of the above theorem shows that the controller is optimal for an index with kernel having a form

$$\bar{L}(\cdot) = \frac{\varepsilon}{\mu}\left\{(\|v_k\|^2_{\bar{R}} + \|w_k\|^2_{\bar{Q}}) - (\|v_k\|^2_{\bar{R}} + \|w_k\|^2_{\bar{Q}})\right\}$$

$$+L(x,u). \tag{24}$$

Again we see that under the conservative design philosophy adopted to cope with plant dynamics or noise uncertainties of optimizing a plant with increased noise and initial state covariances, as in (5), the index for which the controller is optimal is an upper bound for the original performance index used in the design. That is, (23) holds in this more general case also.

We can also see clearly now that for suitably large $\bar{R}, \bar{Q}$, the negative term in the kernels for which the designed controllers are optimal, become negligible and can be ignored in giving an interpretation of the index for which the desired controllers are optimal. The remaining terms can be nonnegative by design and represent control, state and noise costs.

3.3. Linear finite-dimensional controllers

The off-line dynamic programming equation for the index selection of Theorem 3.1 is still infinite-dimensional although not doubly so, as when the information state has an infinite-dimensional representation. The control law is a nonlinear feedback of the information state mean value, see [1]

$$S(\chi,k)=\inf_{u\in U_{k,k}}\bar{E}[S(\chi_{k+1}(\chi_k,u,y_{k+1}),k+1)|\chi_k=\chi] \tag{25}$$

and $S(\chi,T)=\langle q_T(.,\chi),\beta_T\rangle$.

However, in the case that $A(u),B(u),C(u),M(u),N(u)$ are not control dependent and $m(u)$ is quadratic in u (that is $m(u)=u'mu$), then applying the results of [8], the information state mean feedback law is in fact a linear one, and calculated using finite dimensional dynamic programming equations. The resulting control law is given by

$$u_k^{\min}=-(m+B'\tilde{S}_{k+1}B)^{-1}B'\tilde{S}_{k+1}\tilde{A}\mu_k, \tag{26}$$

where μ_k and Σ_k are the mean and variance of the information state respectively, and where

$$\tilde{S}_{k+1}=(S_{k+1}^{-1}-\theta\tilde{R}_k\tilde{R}_k')^{-1},\quad \tilde{A}=A\rho^{-1},$$
$$\tilde{M}=M\rho^{-1},\qquad \rho=I-\theta M\Sigma_k,$$
$$\tilde{R}_k=A\tilde{K}_kC'R^{-1}C\Theta,\quad \Theta=[(C'RC)^{-1}\rho^{-1}\Sigma_k]^{1/2},$$
$$\tilde{K}_k=(\Sigma_k^{-1}+C'R^{-1}C-\theta M)^{-1}.$$

Also, S_k is given by the following backwards recursion:

$$S_k=\tilde{M}+\tilde{A}'S_{k+1}\ (I+Bm^{-1}B'S_{k+1}-\theta\tilde{R}_k\tilde{R}_k'S_{k+1})^{-1}\tilde{A} \tag{27}$$

under the condition that $(I-\theta\tilde{R}_k'S_{k+1}\tilde{R}_k)$ is positive definite for all k.

3.4. Non-gaussian noise

To apply the same concepts as above for plants with non-Gaussian noise, it is appropriate to work with the class of indices (10), where $f(\cdot)$ is able to 'absorb' the non-Gaussian densities involving plant nonlinearities in the same way as $\exp(L)$ 'absorbed' exponential terms derived from Gaussian densities. As a step to see this let us first rewrite (21) as

$$\exp\left\{\frac{\mu}{\varepsilon}\bar{L}(\cdot)\right\}=\frac{\mathcal{N}[v_k,R]\mathcal{N}[w_k,Q]}{\mathcal{N}[v_k+c(x_k),\bar{R}]\mathcal{N}[w_k+a(x_k,u_k),\bar{Q}]}\times\exp\left\{\frac{\mu}{\varepsilon}L(\cdot)\right\}. \tag{28}$$

The generalization of this result to the non-Gaussian version of the nonlinear plant (1), with 'linear' Gaussian monel (2) as before, is

$$f\left(\frac{\mu}{\varepsilon}\bar{L}(\cdot)\right)=\frac{\psi(v_k)\phi(w_k)}{\mathcal{N}[v_k+c(x_k),\bar{R}]\mathcal{N}[w_k+a(x_k,u_K),\bar{Q}]}\times\exp\left\{\frac{\mu}{\varepsilon}L(\cdot)\right\}. \tag{29}$$

Again, further variations on this theme are possible, but not explored further here.

3.5. Risk neutral results

The risk sensitive performance index (7) becomes risk neutral in the limit as $\theta=\mu/\varepsilon$ approaches zero. That is, the exponential term becomes in effect a linear term, so the index is simply $\sum_{i=0}^{M-1}L(x_i,u_i)+\Phi(x_M)$.

Notice that the inverse optimal index (19) under (18) is not a risk neutral index as θ approaches zero. To see this in the simplest manner, take limits as θ approaches zero to both sides of (28), and observe that the Gaussian densities remain exponential and do not become linear. This observation highlights the importance of a risk sensitive optimal control theory even to support the more usual risk neutral theory.

3.6. Small noise limits

In risk sensitive optimal control theory as in [12], small noise limits of the controllers achieve optimal deterministic nonlinear control in the presence of 'worst case' deterministic noise. These can be viewed as nonlinear H_∞ controllers. The key is to work with noise covariances $R<Q$ of order ε in the Gaussian case. Here we do likewise, and also for $\bar{R},\bar{Q}$. Let us generalize the notation as $R^\varepsilon,Q^\varepsilon,\bar{R}^\varepsilon,\bar{Q}^\varepsilon$.

3.7. Continuous-time results

Consider the continuous-time (stochastic) signal model on $t \in 0, T$

$$
\begin{aligned}
\mathrm{d}x_t &= A(u_t)x_t + a(x_t, u_t) + \mathrm{d}w_t, \\
\mathrm{d}y_t &= Cx_t + c(x_t) + \mathrm{d}v_t,
\end{aligned} \tag{30}
$$

where w_t, v_t are Wiener processes with covariance matrices $Q(\cdot), R(\cdot)$, perhaps functions of x_t, u_t. The associated risk sensitivity index is

$$
J(u) = E\left[\exp\left\{\frac{\mu}{\varepsilon}\left(\int_0^T L(x_t, u_t)\,\mathrm{d}t + \Phi(x_T)\right)\right\}\right]. \tag{31}
$$

Define an operator

$$
\mathscr{A} = \sum_i \left[A_i(u_t)x_t + a_i(x_t, u_t)\right]\frac{\delta}{\delta x_i} + \frac{1}{2}\sum_{i,j} Q_{ij}(x_t)\frac{\delta^2}{\delta x_i \delta x_j}. \tag{32}
$$

The information state filter is now given from the modified Zakai equation, see [13]

$$
q(x,t) = q_0(x) + \int_0^t q(x,r)\left[Cx + c(x,r)\right]^* R_r^{-1}\,\mathrm{d}zr + \int_0^t (\mathscr{A}^* q)(x,r)\,\mathrm{d}r + \frac{\mu}{\varepsilon}\int_0^t L(x,r)q(x,r)\,\mathrm{d}r, \tag{33}
$$

where * denotes an adjoint. Again, we see the techniques exploited in Theorem 3.1 can be applied. That is, 'absorb' the nonlinear component of the adjoint operators of (33) into the cost index kernel L so that it becomes (suitably) quadratic.

4. Conclusions

The key new results of the paper are inverse optimal control results for nonlinear control. They give readily interpreted performance indices for which a useful class of controllers is optimal. Also, insightful performance bounds follow as a corollary from these main results. The results apply to controllers designed for suitably simplified plant models and appropriate indices when applied to the actual plant. For simplicity, we have assumed here models for design of the same dimension as the plant equations, and focussed on practical finite dimensional controllers designed for linearized models and risk sensitive indices with quadraticized performance index kernels.

The key technical approach for achieving the results has been to apply information state feedback risk sensitive optimal control theory for general nonlinear (possibly stochastic) systems to a class of performance indices designed to absorb some of the nonlinearities in the plant in such a way as to achieve finite dimensional controllers. The resultant designs are equivalent to those achieved by working with models linearized with respect to the states and possibly the controls.

The results are significant because they expand our understanding of the role of finite dimensional controllers for nonlinear optimal feedback control, and yet exploit the power of the very general optimal nonlinear (possibly stochastic) feedback control theory available now. In this way they flesh out this powerful theory and are a further step to make this theory useful for applications where infinite-dimensional controllers are not practical.

References

[1] L. Aggoun, A. Bensoussan, R.J. Elliott and J.B. Moore, A finite-dimensional quasi-linear risk-sensitive control problem, to appear in: *Systems Control Lett.* (1994).

[2] J.A. Ball and J.W. Helton, Nonlinear H_∞ control for nonlinear systems with output feedback, *IEEE Trans. Automatic Control* **AC-38** (1993) 546–559.

[3] J.S. Baras, Group invariance methods in nonlinear filtering of diffusion processes, *IEEE* (1980) 72–79.

[4] J.S. Baras and M.R. James, Robust output feedback control for discrete-time nonlinear systems: The finite-time case, in: *Proc. 32nd IEEE Conf. on Decision and Control* 51–55.

[5] T. Basar and P. Bernhard, *H_∞ Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach* (Birkhauser, Boston, 1991).

[6] V.E. Beneš and R.J. Elliott, Finite dimensional solutions of a modified Zakai equation, in preparation.

[7] A. Bensoussan and R.J. Elliott, A finite dimensional risk sensitive control problem, to appear in: *SIAM J. Control Optim.*

[8] I.B. Collings, J.B. Moore and M.R. James, An information state approach to linear/risk-sensitive/quadratic/Gaussian control, in: *Proc. 33rd Conf. on Decision and Control*, Florida, USA (1994).

[9] R.J. Elliott, L. Aggoun and J.B. Moore, *Hidden Markov Models: Estimation and Control* (Springer, Berlin, 1994).

[10] R.J. Elliott and J.B. Moore, Discrete time control under a reference measure, *Proc. 12th World Congress IFAC*, Sydney, Australia (1993) 157–160.

[11] M.R. James and J.S. Baras, Robust H_∞ output feedback control for nonlinear systems, *IEEE Trans. Automatic Control* **AC-40** (6) (1995) 1007–1017.

[12] M.R. James, J.S. Baras and R.J. Elliott, Risk-sensitive control and dynamic games for partially observed discrete-time nonlinear systems, *IEEE Trans. on Automatic Control* **39**(4) (1994) 780–792.

[13] M.R. James, J.S. Baras and R.J. Elliott, Output feedback risk-sensitive control and differential games for continuous time nonlinear systems, *Proc. 32nd Conf. on Decision and Control*, San Antonio, TX (1993) 3357–3360.

[14] A. Isidori and A. Astolfi, Disturbance attenuation and H_∞ control via measurement feedback in nonlinear systems, *IEEE Trans. on Automatic Control* **37** (6) (1992) 1283–1293.

[15] D.L. Ocone, J.S. Baras and S.I. Marcus, Explicit filters for diffusions with certain nonlinear drifts, *Stochastics* **8**, 1–16.

[16] A.J. van der Schaft, On a state space approach to nonlinear H_∞ control, *Systems Control Lett.* **16** (1991) 1–8.

[17] A.J. van der Schaft, L_2 gain analysis of nonlinear systems and nonlinear state feedback H_∞ control, *IEEE Trans. on Automatic Control* **AC-37** (6) (1992) 770–784.

[18] A.J. van der Schaft, Nonlinear state space H_∞ control theory, to appear in: H.L. Trentelman and J.C. Willems, Eds., *Perspectives in Control* Series: Progress in Systems Control (Birkhauser, Basel).

[19] P. Whittle, *Risk-sensitive Optimal Control* (Wiley, New York, 1990)

[20] P. Whittle, Risk-sensitive linear/quadratic/Gaussian control, *Adv. Appl. Probab.* **13** (1981) 764–777.