JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 21, No. 5, September–October 1998

Linear-Quadratic Worst-Case Control

Matthew K. Juge* and Arthur E. Bryson Jr.†
Stanford University, Stanford, California 94305

This paper gives an applied view of linear-quadratic worst-case control and relates it to linear-quadratic-Gaussian smoothing. It extends the worst-case control problem formulation to include tracking of desired output histories and nonzero terminal constraints and separates the problem into future, past, and present problems, each of which must satisfy a conjugate-point condition. It includes a finite time horizon example, namely, a helicopter hover-position change in which the disturbance is horizontal wind velocity. Linear-quadratic best-case controllers can be obtained by using positive instead of negative weights in the quadratic performance index.

Nomenclature

A	= system matrix
B_u, B_w	= input injection matrices
C_f, C_c, C_s	= state output matrices
J	= quadratic performance index
$P(t)$	= state covariance matrix
$\hat{P}(t)$	= smoothed covariance
Q_0, Q_f, Q_c, Q_s	= error penalties
$R_u(t)$, $R_w(t)$	= input penalty matrices
$S(t)$	= information matrix
$u(t)$, $w(t)$	= control, disturbance input
v_0, v_f, v_c, v_s	= disturbance/error terms
$x(t)$	= system state
$\hat{x}(t)$	= pessimistic smoothed state estimate
$x(t_0)$	= perturbed initial condition
$x_F(t)$	= state feedforward
x_0	= nominal initial condition
$y_c(t)$	= reference output for tracking
y_f	= desired terminal output
$y_s(t)$	= sensor output
$\lambda_B(t)$	= adjoint

Introduction

DIFFERENTIAL games (minimax problems) in the calculus of variations were first treated by Isaacs.[1] In that same year, Ho et al.[2] considered the special case of linear-quadratic (LQ) differential games, and a chapter on differential games was included in Ref. 3. The viewpoint of regarding the Kalman filter as the forward sweep of a deterministic weighted-least-squares-fit smoothing problem was proposed by Bryson and Frazier.[4] Fraser and Potter[5] gave a forward and backward filter algorithm for such smoothing problems, herein adapted to the linear-quadratic worst-case controller or LQW problem.

Jacobson[6] showed that disturbances could either help or hinder the controls and made the connection to differential games for the latter case with perfect state information, using a linear-exponential-of-quadratic-Gaussian performance index. Speyer et al.[7] extended Ref. 6 to the case with only sensor information. Markov and Reid[8] found a time-varying minimax aircraft controller for worst wind histories, and Chu[9] solved the time-invariant version of the same problem. A different, frequency-domain approach originated with Zames[10] as the H_∞ control design problem. Petersen[11] showed that, with perfect state information, this problem could be solved with a single algebraic Riccati equation. Doyle et al.[12] showed that, with only sensor information, the problem could be reduced to solving two algebraic Riccati equations. Rhee and Speyer[13] considered the finite time horizon LQ differential game and related it to the H_∞ control design problem. Başar and Bernhard[14] provided a comprehensive treatment of the LQ differential game problem in both discrete and continuous time under a number of information structures.

Whittle[15] formulated his risk-sensitive certainty-equivalence principle (RSCEP) and expanded on it in his text on risk-sensitive optimal control.[16] His stochastic, exponential-of-quadratic performance index differs from the deterministic LQW performance index. However, Whittle, Speyer, and others have established that both formulations yield identical controllers. Whittle's risk-averse and risk-seeking controllers are precisely worst-case and best-case controllers.

This paper gives a derivation of the finite time horizon LQW controller, which is similar to Whittle's derivation[16] of the RSCEP. However, it uses the LQW quadratic performance index (QPI) and the smoothing concept of Bryson and Frazier.[4] The derivation extends the controller of Rhee and Speyer[13] to include tracking of desired outputs and to handle nonzero terminal constraints. LQW controller synthesis can be done using standard linear-quadratic-Gaussian (LQG) computational algorithms, with the simple modification of allowing indefinite weighting matrices in the quadratic performance indices and combining forward and backward sweeps.

The choice of the weighting matrices in the QPIs requires knowledge of the system to be controlled and engineering judgment about the objectives of the controller design. The authors believe there is still much to be learned about how to do this before LQW is accepted as an engineering design tool. To this end, this paper attempts to provide an intuitive link between the aforementioned H_∞ literature and the existing body of LQG (H_2) knowledge.

The derivation consists of three parts.

1) The first part is a future problem of finding control and disturbance feedback strategies with an arbitrary initial state at the current time $t = t_1$. This is a differential game because the disturbances maximize the output errors that the controls are minimizing. The controls and the disturbances are bounded by integral-quadratic penalties. The controller is limited to sensor information.

2) The second part is a past problem of estimating the initial conditions and the past disturbance histories using the past sensor measurements with the final state at $t = t_1$ open. The past control history is known and fixed. The smoothing performance index assumes that the disturbance $w(t)$ tries to maximize the tracking error and is bounded by an integral-quadratic penalty.

3) The third part is a present problem of joining the future and past solutions at the current time $t = t_1$. The future-influenced current state estimate is determined by a simple optimization of the difference of two quadratic forms involving the current state.

Riccati equations with indefinite weighting matrices may tend to infinity in a finite time; a point at which this occurs is called a conjugate point in the classical calculus of variations. Conjugate

Presented as Paper 96-3810 at the AIAA Guidance, Navigation, and Control Conference, San Diego, CA, July 29–31, 1996; received May 19, 1997; revision received March 19, 1998; accepted for publication April 2, 1998.

*Doctoral Candidate, Department of Aeronautics and Astronautics; currently Engineering Specialist, Guidance, Control, and Performance Department, Raytheon Missile Systems Co., P.O. Box 11337, Building 805, Station M4, Tucson, AZ 85734-1337. Senior Member AIAA.

†Deptartment of Aeronautics and Astronautics, Durand Building. Honorary Fellow AIAA.

points may appear in the solutions to problems 1, 2, or 3 if the negative weights in the QPI are chosen too small relative to the positive weights. These weights are too small if the increase in cost due to a disturbance can ever exceed the penalty cost imposed on that disturbance, in which case the disturbance increases without bound and yields infinite cost. Limebeer et al.[17] relate this aspect of conjugate point behavior to the H_∞ limit γ^2. The requirement that no conjugate points occur along the solution path is equivalent to the conditions for a differential game saddle-point solution to exist, as described by Juge.[18]

The results are specialized to the infinite time horizon problem for time-invariant systems. We briefly discuss best-case or risk-seeking controllers. Position control of an OH-6A helicopter demonstrates implementation of the controllers described herein.

Worst-Case LQ Controller

The worst-case LQ controller problem may be posed as finding the control vector $u(t)$ to minimize and the disturbances to dynamics, measurements, and initial conditions $w(t)$, $v_s(t)$, $x(t_0)$ to maximize

$$J = \frac{1}{2}\left[v_f^T Q_f v_f - v_0^T Q_0 v_0 + \int_{t_0}^{t_f} \left(v_c^T Q_c v_c - v_s^T Q_s v_s + u^T R_u u - w^T R_w w\right) \mathrm{d}t\right] \tag{1}$$

subject to

$$\dot{x} = Ax + B_u u + B_w w \tag{2}$$

where

$$v_f = C_f x(t_f) - y_f \tag{3}$$

$$v_0 = x(t_0) - x_0 \tag{4}$$

$$v_c(t) = C_c x(t) - y_c(t) \tag{5}$$

$$v_s(t) = C_s x(t) - y_s(t) \tag{6}$$

The desired terminal conditions y_f, the nominal initial conditions x_0, and the desired tracking output $y_c(t)$ are specified. At any given time t_1, the sensor data $y_s(t)$ and control history $u(t)$ are known and fixed in the past; $y_s(t)$ in the future is not yet known to the controller.

The four terms given by Eqs. (3–6) and appearing quadratically in Eq. (1) are the terminal error v_f, the perturbation to initial conditions v_0 introduced by the disturbance's choice of $x(t_0)$ away from the nominal value x_0, the tracking error $v_c(t)$, and the sensor noise $v_s(t)$.

The control $u(t)$ and the disturbances $w(t)$ and $v_s(t)$ are restricted to the class of square-integrable functions, whereas the perturbation to initial conditions v_0 is restricted to be finitely bounded. The control is further restricted to be causal, i.e., independent of future values of $v_s(t)$.

Separation into Future and Past Problems

Because $u(t)$ is fixed in the past and the future sensor error $v_s(t)$ cannot influence a causal controller, Eq. (1) can be written, except for constant terms, as the difference between a future QPI and a past QPI:

$$J = J_{\text{Future}} - J_{\text{Past}} \tag{7}$$

where

$$J_{\text{Future}} = \frac{1}{2}\left[v_f^T Q_f v_f + \int_{t_1}^{t_f} \left(v_c^T Q_c v_c + u^T R_u u - w^T R_w w\right) \mathrm{d}t\right] \tag{8}$$

$$J_{\text{Past}} = \frac{1}{2}\left[v_0^T Q_0 v_0 + \int_{t_0}^{t_1} \left(v_s^T Q_s v_s - v_c^T Q_c v_c + w^T R_w w\right) \mathrm{d}t\right] \tag{9}$$

and t_1 = the present time. The control u does not appear in the past cost, as it is a now-determined, present quantity; in a symmetric fashion, the sensor noise v_s does not appear in the future cost, as it cannot affect a causal controller in the present. The future problem is a generalized LQ terminal control problem with an arbitrary $x(t_1)$. The past problem is a generalized LQ filtering problem, also with $x(t_1)$ open. These problems are exact duals of each other. Each problem may be solved independently; the results are then combined by sweeping the initial conditions forward for the past problem and the final conditions backward for the future problem and maximizing to find the worst-case current state $x(t_1)$.

Future Problem

The future problem is to find the control vector history $u(t)$ to minimize and the disturbance vector history $w(t)$ to maximize Eq. (8) subject to Eq. (2) for arbitrary $x(t_1)$. This problem can be reduced in appearance to a standard LQ terminal control problem by augmenting the control vector to include the disturbance vector as an unfriendly control whose weighting matrix is negative definite, i. e., find $\bar{u}(t)$ to determine an extremum of

$$J_{\text{Future}} = \frac{1}{2}\left[v_f^T Q_f v_f + \int_{t_1}^{t_f} \left(v_c^T Q_c v_c + \bar{u}^T \bar{R}\bar{u}\right) \mathrm{d}t\right] \tag{10}$$

subject to

$$\dot{x} = Ax + \bar{B}\bar{u} \tag{11}$$

where

$$\bar{u} \triangleq \begin{bmatrix} u \\ w \end{bmatrix}, \qquad \bar{B} \triangleq [B_u \quad B_w], \qquad \bar{R} \triangleq \begin{bmatrix} R_u & 0 \\ 0 & -R_w \end{bmatrix} \tag{12}$$

Although this is a minimax problem, the first-order necessary conditions for a stationary solution do not distinguish between minima, maxima, and saddle points. The extremal solution is the well-known full-state LQ terminal controller; it consists of a feedforward control $\bar{u}_f(t)$ plus a feedback control $-\bar{K}(t)x$:

$$\bar{u}^* = \bar{u}_f(t) - \bar{K}(t)x \tag{13}$$

where

$$\bar{u}_f = \bar{R}^{-1}\bar{B}^T\lambda_B(t), \qquad \bar{K} = \bar{R}^{-1}\bar{B}^T S(t) \tag{14}$$

and

$$-\dot{S} = SA + A^T S + C_c^T Q_c C_c - S\bar{B}\bar{R}^{-1}\bar{B}^T S \qquad S(t_f) = C_f^T Q_f C_f \tag{15}$$

$$-\dot{\lambda}_B = [A - \bar{B}\bar{K}]^T\lambda_B + C_c^T Q_c y_c, \qquad \lambda_B(t_f) = C_f^T Q_f y_f \tag{16}$$

Applying Eq. (12) to eliminate the overbarred variables from Eqs. (13–16) yields

$$\begin{bmatrix} u^* \\ w^* \end{bmatrix} = \begin{bmatrix} u_f(t) - K_u(t)x \\ w_f(t) - K_w(t)x \end{bmatrix} \tag{17}$$

$$\begin{bmatrix} u_f \\ w_f \end{bmatrix} = \begin{bmatrix} R_u^{-1}B_u^T \\ -R_w^{-1}B_w^T \end{bmatrix}\lambda_B(t) \tag{18}$$

$$\begin{bmatrix} K_u \\ K_w \end{bmatrix} = \begin{bmatrix} R_u^{-1}B_u^T \\ -R_w^{-1}B_w^T \end{bmatrix} S(t) \tag{19}$$

$$\bar{B}\bar{R}^{-1}\bar{B}^T = B_u R_u^{-1}B_u^T - B_w R_w^{-1}B_w^T \tag{20}$$

$$\bar{B}\bar{K} = B_u K_u - B_w K_w \tag{21}$$

On pages 182 and 183 of Ref. 3, McReynolds showed how to complete the square to prove minimality for the case where $\bar{R}$ is positive definite. This procedure also proves the minimax or saddle-point nature of the preceding solution for $\bar{R}$ indefinite as shown by

Juge[18]; deviations from the extremal control and disturbance yield an increment in cost

$$\Delta J_{\text{Future}} = \frac{1}{2}\left[v_1^T S(t_1) v_1 + \int_{t_1}^{t_f} (u - u^*)^T R_u (u - u^*)\,\mathrm{d}t - \int_{t_1}^{t_f} (w - w^*)^T R_w (w - w^*)\,\mathrm{d}t\right] \tag{22}$$

where

$$v_1 \triangleq x(t_1) - x_B(t_1), \qquad x_B(t) = S^{-1}(t)\lambda_B(t) \tag{23}$$

The solution is a saddle point for $S \geq 0$, because the cost will increase if the controls do not use the optimal feedback $u^*(t)$ and will decrease if the disturbances are not $w^*(t)$. In fact, a conjugate point may occur in solving Eq. (10) if R_w is chosen too small relative to R_u; i.e., $S(t)$ may tend to infinity in a finite time. For a realistic solution, the problem should be iterated (R_w reduced) until the magnitudes of the disturbances and controls are reasonable. The term Q_f is positive definite, and so $S(t_f) = C_f^T Q_f C_f$ is at least positive semidefinite. The term $S(t)$ cannot become indefinite without first becoming infinite, which corresponds to a conjugate point.

Past Problem

This problem is a dual minimax problem between the process and measurement disturbances $w(t)$ and $v_s(t)$ on the one hand and the tracking error $v_c(t)$ on the other. The term J_{Past} does not involve $u(t)$ because the past control is fixed. Positive penalties on sensor noise $v_s(t)$ and process disturbances $w(t)$ and a negative penalty on tracking error $v_c(t)$ imply that disturbances act to reduce J_{Past} by increasing the tracking error. [Recall the negative sign on J_{Past} in Eq. (7); the disturbance here seeks to minimize this negated cost to maximize overall cost.]

This problem can be reduced in appearance to a linear-quadratic smoother problem by augmenting the measurement vector $y_s(t)$ to include the tracking vector $y_c(t)$ with a negative-definite weighting matrix. The term $v_s(t)$ is determined by the known history $y_s(t)$ and the state history $x(t)$, which, in turn, is determined from Eq. (2) by $w(t)$, $x(t_0)$, and the known history $u(t)$. Hence the problem becomes one of finding $w(t)$ and $x(t_0)$ to minimize:

$$J_{\text{Past}} = \frac{1}{2} v_0^T Q_0 v_0 + \frac{1}{2}\int_{t_0}^{t_1} \left(\bar{v}^T \bar{Q} \bar{v} + w^T R_w w\right)\mathrm{d}t \tag{24}$$

where

$$\bar{y} \triangleq \begin{bmatrix} y_s \\ y_c \end{bmatrix}, \qquad \bar{C} \triangleq \begin{bmatrix} C_s \\ C_c \end{bmatrix}, \qquad \bar{v} \triangleq \bar{C}x - \bar{y}, \qquad \bar{Q} \triangleq \begin{bmatrix} Q_s & 0 \\ 0 & -Q_c \end{bmatrix} \tag{25}$$

This is a generalized LQG problem because the weighting matrix $\bar{Q}$ on the in-flight outputs $\bar{y}(t)$ is not positive definite. A first-order extremal solution to this problem is well known; the solution is a generalization of the Kalman–Bucy filter:

$$\dot{x}_F = Ax_F + B_u u + \bar{L}(t)(\bar{y} - \bar{C}x_F), \qquad x_F(t_0) = x_0 \tag{26}$$

where

$$\bar{L}(t) = P(t)\bar{C}^T\bar{Q} \tag{27}$$

and

$$\dot{P} = AP + PA^T + B_w^T R_w^{-1} B_w - P\bar{C}^T\bar{Q}\bar{C}P, \qquad P(t_0) = Q_0^{-1} \tag{28}$$

This forward filter sweeps the nominal initial conditions to equivalent conditions at $t = t_1$; any deviation of the disturbances from their optimal values that produces an achieved current state $x(t_1)$ other than the worst-case state $x_F(t_1)$ results in a cost increment

$$\Delta J_{\text{Past}} = \tfrac{1}{2}[x(t_1) - x_F(t_1)]^T P^{-1}(t_1)[x(t_1) - x_F(t_1)] \tag{29}$$

The saddle-point condition $P > 0$ ensures that no conjugate points occur in the forward sweep.

Equation (26) differs from the estimator equation typically found in the H_∞ literature.[12,13] The conventional approach is to propagate an estimate of the worst-case value of the actual state, which requires both a bias term to reflect the action of the estimated worst-case disturbance as well as a modification to the preceding estimator gain L that accounts for future effects. The approach taken here simply propagates the nominal initial conditions forward to the present time using an unbiased estimator and feeding back sensor innovations and estimated tracking error. The worst-case state is estimated at the present time by combining this nominal estimate with the backward sweep from the future, as seen in the next section.

Combining Future and Past Problems

The optimal swept-forward solution of the past problem may now be combined with the optimal swept-back solution of the future problem in a manner reminiscent of the forward and backward filter smoother.[5] Combining cost variations from the future, Eq. (22), and from the past, Eq. (29), in Eq. (7) leads to a present variation

$$\Delta J = \tfrac{1}{2}[x(t_1) - x_B(t_1)]^T S(t_1)[x(t_1) - x_B(t_1)] - \tfrac{1}{2}[x(t_1) - x_F(t_1)]^T P^{-1}(t_1)[x(t_1) - x_F(t_1)] \tag{30}$$

The term ΔJ is to be maximized with respect to the state $x(t_1)$; that is, the worst $x(t_1)$ is found that is consistent with past sensor data and potential future disturbances. The solution is straightforward when ΔJ is written as

$$\Delta J = -\tfrac{1}{2}[x(t_1) - \hat{x}(t_1)]^T \hat{P}^{-1}(t_1)[x(t_1) - \hat{x}(t_1)] \tag{31}$$

where

$$\hat{x}(t_1) = \hat{P}(t_1)\left[P^{-1}(t_1)x_F(t_1) - S(t_1)x_B(t_1)\right] \tag{32}$$

$$\hat{P}^{-1}(t_1) = P^{-1}(t_1) - S(t_1) \tag{33}$$

For a minimum to occur, the future-influenced weighting matrix $\hat{P}(t_1)$ must be positive definite. A conjugate point may occur in the forward integration of Eq. (28) if Q_c is chosen too large relative to Q_s; i.e., $P(t)$ may tend to infinity in a finite time. For a realistic solution, the problem should be iterated (Q_c decreased) until the magnitudes of the disturbances and outputs are reasonable.

The saddle-point worst-case state is then $\hat{x}(t_1)$; any deviation of the actual state $x(t_1)$ from $\hat{x}(t_1)$ reduces the cost, counter to the interests of a maximizing disturbance.

The term $\hat{x}(t_1)$ is also the worst-case future-influenced state estimate, which results from the past actions of $w(t)$ and $x(t_0)$ and the predicted actions of $w(t)$ in the future. Interpreting $P^{-1}(t_1)$ as a forward information matrix and $S(t_1)$ as a negative information matrix, the future-influenced weighting matrix $\hat{P}(t_1)$ is increased by the negative information about the predicted future, i.e., $\hat{P}(t_1) > P(t_1)$. This results from the conservative (or pessimistic) assumption of worst-case disturbances both in the past and in the future. The form of this solution is a straightforward generalization of the Fraser–Potter smoother,[5] which consists of forward and backward information filters. Equation (31) shows the saddle-point nature of the maximizing solution with respect to $x(t_1)$ for $\hat{P} > 0$; any deviation of $x(t_1)$ from $\hat{x}(t_1)$ will decrease J.

The estimator detects deviations of the disturbances from their worst values, allowing the controller to take advantage of any such deviations to improve performance. Only if the disturbances are the worst expected and the controller uses its optimal strategy will the actual state take on its minimax value $\hat{x}$, the future-influenced state estimate. Any deviations of the disturbance from optimal costs the disturbance more than exciting the estimator error mode is worth; this is a direct consequence of the saddle-point conditions and is discussed at greater length in Chapter 6 of Ref. 18.

In the standard LQG problem, the $w(t)$ disturbance term does not appear in the future QPI (8), and the $v_c(t)$ tracking term does not appear in the past QPI (9), so that the future and past problems are decoupled, and

$$\hat{x}(t_1) = x_F(t_1), \qquad \hat{P}(t_1) = P(t_1) \tag{34}$$

confirming the well-known certainty-equivalence principle.

The final result is that the optimal control at $t = t_1$ is given by

$$u(t_1) = u_f(t_1) - K_u(t_1)\hat{x}(t_1) \tag{35}$$

or

$$u(t_1) = R_u^{-1}B_u^T\left\{\lambda_B(t_1) - S(t_1)[I - P(t_1)S(t_1)]^{-1}\right\} \times [x_F(t_1) - P(t_1)\lambda_B(t_1)] \tag{36}$$

The cost when both the controller and the disturber play their optimal strategies is the value V of the differential game; after tedious but straightforward manipulation, substitution of the optimal solutions into Eq. (1) yields

$$V = \tfrac{1}{2}x_0^T[S^{-1}(0) - P(0)]^{-1}x_0 \tag{37}$$

where x_0 again represents the nominal initial conditions and S and P are the solutions to the future and past Riccati equations.

In the regulator problem, $y_f = 0$ and $y_c(t) = 0$, so that $\lambda_B(t) = 0$, no feedforward is present, and Eq. (36) simplifies to

$$u(t_1) = -K_u(t)[I - P(t_1)S(t_1)]^{-1}x_F(t_1) \tag{38}$$

Example: Helicopter Hover-Position Change in Gusty Winds

Here controller logic is synthesized to change the position of a hovering helicopter in the presence of gusty winds, as in the example used in Ref. 19. The control is cyclic pitch of the rotor, and the disturbance is the horizontal wind velocity, the value of which is unknown to the controller. Two cases are considered: one with a position sensor only and the other with position and pitch angle sensors.

The term R_w is adjusted to give a worst wind gust history that has a maximum velocity of 15 ft/s, and R_u is adjusted to give a feasible maximum cyclic stick deflection. The term Q_f is chosen so that the final conditions are reasonably close to equilibrium in the new hover position. The helicopter model and the QPI parameters are given in the Appendix. Motion of the controlled helicopter is simulated with the worst wind history and from the worst initial conditions.

Figure 1 shows the position, ground speed, and attitude response of the controlled helicopter to a position-change command of 10 ft, with the worst wind gust velocity history and the worst initial condition perturbations. Figure 2 shows the cyclic control and wind disturbance histories. Positive eigenvalues of S, P, $\hat{P} = [P^{-1} - S]^{-1}$ were checked to ensure that the solution is valid, i.e., a saddle point.

Figures 1 and 2 are the same for both one or two sensors, but satisfaction of the $\hat{P} > 0$ condition requires that the weight on sensor noise with position measurement alone has to be 100 times

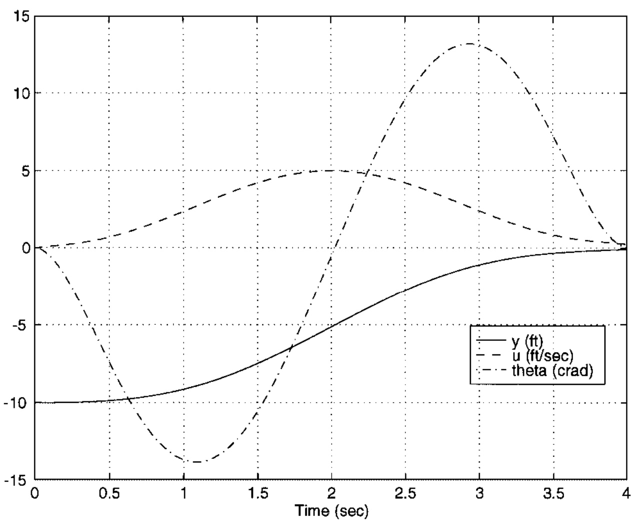

Fig. 1 Helicopter hover-position change with worst wind gust and initial condition perturbations: position *y*, ground velocity *u*, and pitch attitude θ vs time.

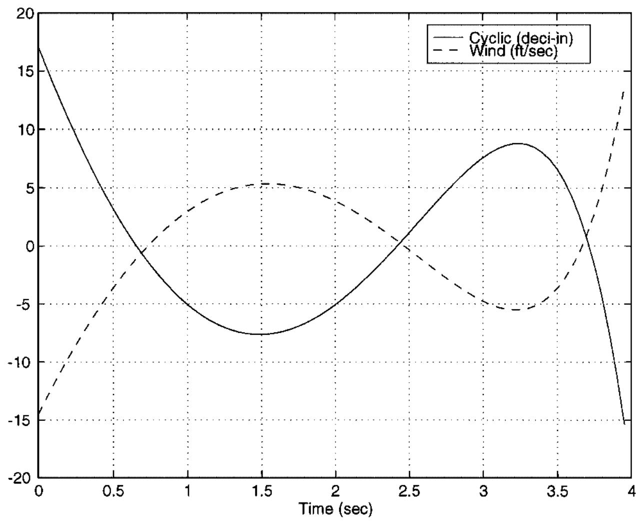

Fig. 2 Helicopter hover-position change with worst wind gust and initial condition perturbations: cyclic pitch control δ and wind disturbance *w* vs time.

larger than with both position and pitch angle available. In other words, estimation with a position sensor alone requires a sensor accuracy two orders of magnitude finer than with both position and pitch sensors.

Time-Invariant LQW Controller with Sensor Feedback

If the plant is time invariant (TI) and both t_0 and t_f become distant horizons, then $S(t)$ and $P(t)$ may become constant matrices, so that the feedback gains K_u, K_w, L_s, and L_c also become constant. This is called a suboptimal H_∞ controller in the literature. The positivity conditions on S, P, and $\hat{P}$ needed to ensure that no conjugate points arise correspond to the conditions given by Doyle[12] (Theorem 3) for an admissible suboptimal H_∞ controller. Conjugate points arise as the penalty weights on disturbances are made smaller; this corresponds to the reduction of the multiplier γ^2 below the H_∞ limit in the notation employed by Doyle. If a realistic controller is desired, the magnitudes of the negative weights need to remain above this limit.

TI LQW controllers can be synthesized using standard LQG software, e.g., the MATLAB LQR and LQE (or DLQR and DLQE) commands, when the following modifications are employed.

1) Remove the checks on positive definiteness of the input weighting matrices.

2) Postmultiply the state-feedback gain matrix (or premultiply the estimator gains) by the amplification factor $[I - PS]^{-1}$.

3) Include a check for positive definiteness of the output matrices S, P, and $\hat{P} = [P^{-1} - S]^{-1}$; these checks replace the safeguards removed in step 1.

Example: Hover-Position-Hold Autopilot for a Helicopter

This is a TI controller version of the previous example. The inflight controlled output is hover position, and the sensed outputs are position and pitch angle. To test the controller, the system response is simulated with the worst wind history and the worst perturbation to initial conditions from a nominal initial condition. The nominal initial conditions x_0 are chosen proportional to the eigenvector corresponding to the largest eigenvalue of the closed-loop cost matrix $(S^{-1} - P)^{-1}$ seen in Eq. (37) for the value of the differential game; this choice of nominal initial condition yields the largest possible excitation of the system for a unit initial condition vector. The worst initial condition as a perturbation from nominal is then $x_w(0) = \hat{P}P^{-1}x_0$; the estimator feedforward x_F is initialized to the nominal initial state. As a consequence of this initialization, the estimator accurately determines the saddle-point worst state, and the state estimate is perfect as long as the disturbance plays its optimal strategy. The Appendix gives the numerical values used in the synthesis and the simulation.

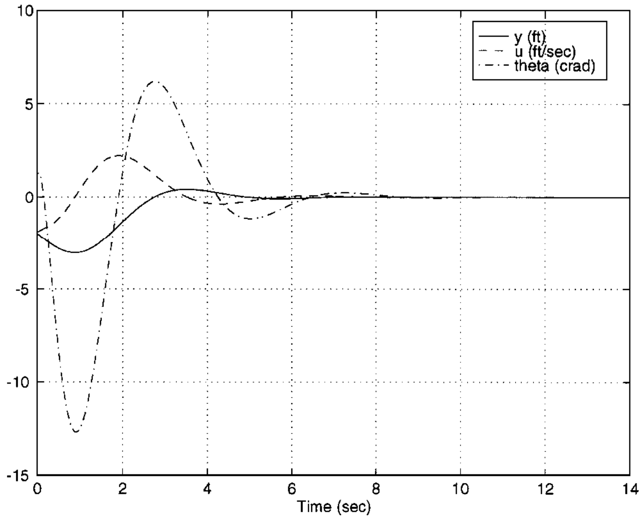

Fig. 3 Response of helicopter with a constant gain LQW position-hold autopilot to worst wind history and worst initial conditions: position y, ground velocity u, and pitch attitude θ vs time.

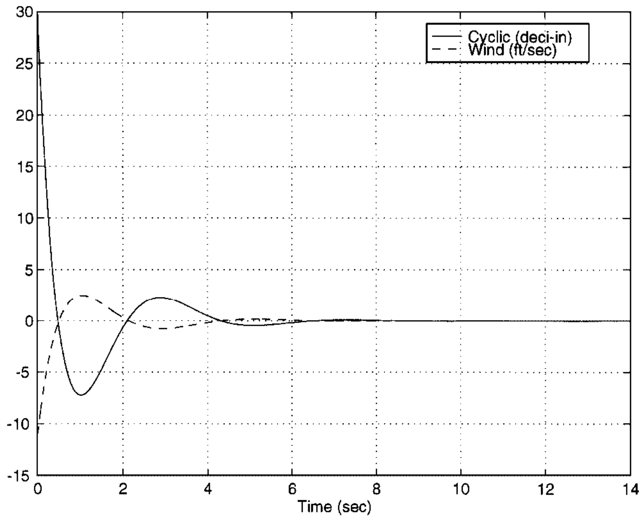

Fig. 4 Response of helicopter with a constant gain LQW position-hold autopilot to worst wind history and worst initial conditions: cyclic pitch control δ and wind disturbance w vs time.

Figures 3 and 4 show the controlled response for the case where initial state vector magnitude was chosen so that the initial position error was 2 ft. The term R_w is chosen to produce a maximum magnitude of the wind velocity approximately equal to 10 ft/s. The resulting initial ground velocity is 1.9 ft/s. The worst wind velocity is generally in the same direction as the ground velocity, and its variation with time is inside the controller bandwidth.

Best-Case Controllers

Linear-quadratic best-case controllers can also be synthesized using LQG codes with only minor modifications. Such controllers involve positive weighting matrices where negative weighting matrices are used in the LQW controller. Such controllers correspond to the risk-seeking controllers described by Whittle[16] because they are based on the optimistic prediction that the disturbances will help control the system in a team effort with the actual controls. This is, of course, not a conservative assumption and is not of much interest to control designers. As an extreme example, wind alone can change the position of the helicopter from the first example, as shown in Figs. 5 and 6. The wind velocity required is very large, on the order of 50 ft/s. Nominal initial conditions are used in this simulation instead of the best initial conditions.

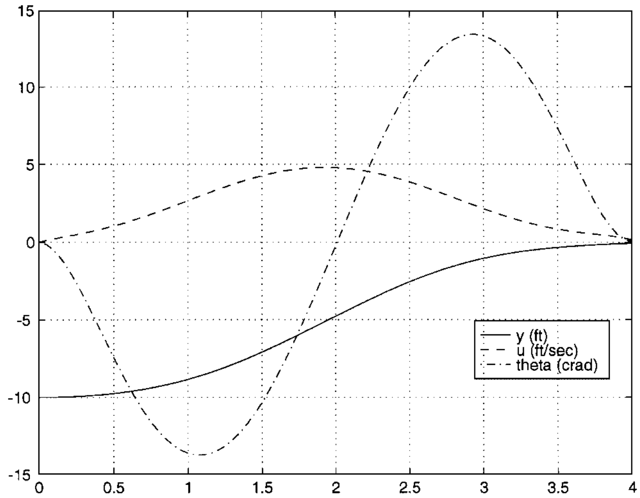

Fig. 5 Extreme example of best-case control; wind variation alone changes the position of a helicopter by 10 ft: position y, ground velocity u, and pitch attitude θ vs time.

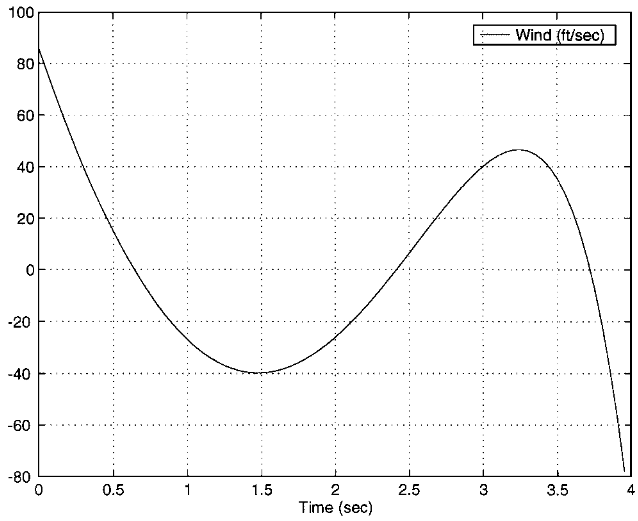

Fig. 6 Extreme example of best-case control: wind velocity w vs time.

Conclusions

Linear-quadratic worst-case controllers can be synthesized using linear-quadratic-Gaussian codes with only minor modifications. The primary differences between LQG and LQW controllers are the following.

1) Worst disturbances and worst deviations of the initial conditions, both with quadratic bounds, are assumed in the synthesis of the controller.

2) The disturbances are assumed to increase the in-flight output errors and to anticipate the future in such a way as to increase the terminal output errors. This requires a backward sweep update of the generalized Kalman filter, here termed a future-influenced state estimate. The algorithm is a generalization of an LQG smoothing algorithm.

3) Conjugate points may arise in the solution of the pair of mixed-definite Riccati equations corresponding to the future and past problems or in the combination of the Riccati solutions in the present problem. Such conjugate points indicate that disturbances exist that can yield infinite cost; a tighter quadratic bound (higher penalty) must be placed on the disturbances for a finite saddle-point solution to exist.

Linear-quadratic best-case controllers can also be synthesized using LQG design software. Such controllers use positive weighting matrices on the disturbance and thus reflect an optimistic best-case solution.

Appendix: Data for Helicopter Hover Examples

The states (in order) are u = horizontal velocity, q = pitch rate, θ = pitch angle, and y = horizontal distance from a reference point. The units are feet, seconds, and centi-radians (0.01 rad). The control is longitudinal cyclic stick in deci-inches (0.1 in.), and the disturbance is w = horizontal wind velocity. The continuous system matrices are

$$A = \begin{bmatrix} -0.0257 & 0.013 & -0.322 & 0 \\ 1.26 & -1.765 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$$

$$B_u = \begin{bmatrix} 0.086 \\ -7.41 \\ 0 \\ 0 \end{bmatrix}, \qquad B_w = \begin{bmatrix} 0.0257 \\ -1.26 \\ 0 \\ 0 \end{bmatrix}$$

The sensed outputs were position and pitch angle

$$C_s = \begin{bmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

For the finite time cases (Figs. 1–6), the terminal output constaints are $y_f = [0\ \ 0\ \ 0\ \ 0]^T$ where $C_f = I_4$, $Q_f = 0.1 I_4$, and there are no in-flight outputs, so that $C_c = [0\ \ 0\ \ 0\ \ 0]$, and $Q_c = 0$. The nominal initial state and weighting matrix for perturbations are

$$x_0 = [0 \quad 0 \quad 0 \quad -10]^T, \qquad Q_0 = I_4$$

The final time is $t_f = 4$ s. The examples are calculated using discrete algorithms with $N = 80$ steps. The discrete weighting matrices are

$$R_u = (1/20^2)/N, \qquad R_w = (1/42^2)/N$$

where the critical value of $R_w = (1/60^2)/N$, where a conjugate point first appears in the computation of S. The discrete weighting matrix for the sensed measurements is

$$Q_s = \begin{bmatrix} 1/0.1^2 & 0 \\ 0 & 1/0.1^2 \end{bmatrix} \Big/ N$$

For the case with only position sensed,

$$C_s = [0 \quad 0 \quad 0 \quad 1], \qquad Q_s = (1/0.001^2)/N$$

The time-invariant case is calculated using the continuous codes LQE and LQR in MATLAB. The in-flight controlled output is position

$$C_c = [0 \quad 0 \quad 0 \quad 1], \qquad Q_c = 1$$

with weighting matrices

$$R_u = 1/4^2, \qquad R_w = 1/6^2$$

The magnitude of the worst initial state vector is chosen to be 4.91 to give an initial position error of −2 ft.

Acknowledgments

The authors built on research performed by Raymond Mills, Laurent El Ghaoui, Alain Carrier, Solo Hermelin, and Peter Chu, who were supported by the Air Force Office of Sponsored Research, France, Lockheed, Israel, and NASA, respectively. The first author gratefully acknowledges support from a National Science Foundation Fellowship and a Hughes Doctoral Fellowship.

References

[1] Isaacs, R., *Differential Games*, Wiley, New York, 1965, Chap. 3.

[2] Ho, Y. C., Bryson, A. E., and Baron, S., "Differential Games and Optimal Pursuit-Evasion Strategies," *IEEE Transactions on Automatic Control*, Vol. 10, No. 10, 1965, pp. 385–389.

[3] Bryson, A. E., and Ho, Y. C., *Applied Optimal Control*, Hemisphere, Washington, DC, 1975, Chap. 5.

[4] Bryson, A. E., and Frazier, M., "Smoothing for Linear and Nonlinear Dynamic Systems," Proceedings of the Optimum Systems Synthesis Conf., U.S. Air Force TR ASD-TDR-63-119, Feb. 1963.

[5] Fraser, D. C., and Potter, J. E., "The Optimum Linear Smoother as a Combination of Two Optimum Linear Filters," *IEEE Transactions on Automatic Control*, Vol. AC-14, No. 4, 1969, pp. 387–390.

[6] Jacobson, D. H., "Optimal Stochastic Linear Systems with Exponential Performance Criteria and Their Relation to Deterministic Differential Games," *IEEE Transactions on Automatic Control*, Vol. AC-18, No. 2, 1973, pp. 124–131.

[7] Speyer, J. L., Deyst, J., and Jacobson, D. H., "Optimization of Stochastic Linear Systems with Additive Measurement and Process Noise Using Exponential Performance Criteria," *IEEE Transactions on Automatic Control*, Vol. AC-19, No. 4, 1974, pp. 358–366.

[8] Markov, A. B., and Reid, L. D., "Conflict of Interest Modeling in Aircraft Response Study," *Journal of Aircraft*, Vol. 19, No. 12, 1982, pp. 1025–1031.

[9] Chu, P. Y., "Control of Aircraft Under Severe Wind Shear Conditions," Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, Dec. 1986.

[10] Zames, G., "Feedback and Optimal Sensitivity: Model Reference Transformations, Multiplicative Seminorms and Approximate Inverses," *IEEE Transactions on Automatic Control*, Vol. AC-26, No. 2, 1981, pp. 301–320.

[11] Petersen, I. R., "Disturbance Attenuation and H-Infinity Optimization: A Design Method Based on the Algebraic Riccati Equation," *IEEE Transactions on Automatic Control*, Vol. AC-32, No. 5, 1987, pp. 427–429.

[12] Doyle, J. C., Glover, K., Khargonekar, P. P., and Francis, B. A., "State-Space Solutions to Standard H_2 and H_∞ Control Problems," *IEEE Transactions on Automatic Control*, Vol. 34, No. 8, 1989, pp. 831–847.

[13] Rhee, I., and Speyer, J. L., "A Game Theoretic Controller and Its Relationship to H_∞ and Linear-Exponential-Gaussian Synthesis," *IEEE Transactions on Automatic Control*, Vol. AC-36, No. 9, 1991, pp. 1021–1032.

[14] Başar, T., and Bernhard, P., *H_∞-Optimal Control and Related Minimax Design Problems*, Birkhäuser, Boston, MA, 1991, Chaps. 5, 6.

[15] Whittle, P., "Risk-Sensitive Linear-Quadratic-Gaussian Control," *Advances in Applied Probability*, Vol. 13, No. 4, 1981, pp. 764–777.

[16] Whittle, P., *Risk-Sensitive Optimal Control*, Wiley, New York, 1990, Chap. 5.

[17] Limebeer, D. J. N., Anderson, B. D. O., Khargonekar, P. P., and Green, M., "A Game Theoretic Approach to H_∞ Control for Time-Varying Systems," *SIAM Journal on Control and Optimization*, Vol. 30, No. 2, 1992, pp. 262–283.

[18] Juge, M. K., "A Differential Game, Time-Domain Interpretation of the H_∞ Design Problem," Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, March 1996.

[19] Bryson, A. E., and Carrier, A., "A Comparison of Control Synthesis Using Differential Games vs. LQR," *Proceedings of the AIAA Guidance, Navigation, and Control Conference* (Boston, MA), AIAA, Washington, DC, Aug. 1989, pp. 1422–1427.