JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 15, No. 5, September–October 1992

Parameter-Robust Control Design Using a Minimax Method

Raymond A. Mills* and Arthur E. Bryson Jr.†
Stanford University, Stanford, California 94305

A controller design method is presented that gives the best linear-quadratic-Gaussian closed-loop performance over a set of worst plant parameter changes. The design algorithm combines a multiplant optimal design code, SANDY, with a new worst parameter algorithm that uses a quadratic norm on parameter changes. The minimax algorithm is unique in the way it weights worst plants to expand the stable region in the parameter space. The method is applied to a two-mass/spring American Control Conference "benchmark" problem. A minimax controller is first designed for the case where the spring constant alone is uncertain. Next, several minimax controllers, including a reduced-order design, are synthesized for the benchmark problem where both masses and the spring constant are uncertain. The results show that minimax control provides near-optimal nominal performance with significant robustness and parameter margin improvements.

Introduction

LINEAR-quadratic-Gaussian (LQG) and $\mathcal{H}_\infty$ controller designs for plants containing lightly damped oscillatory modes are, in general, not robust to plant parameter changes, since they tend to produce finely tuned notch compensators. This paper presents a minimax design method that optimizes closed-loop performance over a range of worst parameter changes. The performance measure used here is the LQG criterion but, in principle, the method can also be applied to the $\mathcal{H}_\infty$ criterion.

Doyle[1] has suggested a method for analyzing the robustness to parameter changes of closed-loop systems, which he calls μ-analysis. The σ-analysis algorithm used here for determining worst parameter changes was given in Ref. 2 for a linear-quadratic-worst-initial-condition (LQW) criterion. It differs from μ-analysis in two ways. First, it uses a quadratic norm for the parameter changes instead of an infinity norm, since parameter variations tend to be Gaussian. Second, it finds the worst parameter changes for a given norm σ on these changes; as σ is increased, a critical value $\sigma_{\max}$ is reached where the closed-loop system is unstable. This value is nearly equal to $1/\mu$, since $1/\mu$ is the half-side of the largest hypercube in the normalized parameter-change space for which the closed-loop system remains stable, whereas $\sigma_{\max}$ is the radius of the largest hypersphere.

The σ-analysis is iterative but converges to a local extremum in only a few iterations. The problem is nonconvex so this local extremum may not, of course, be the global extremum. As yet there do not appear to be efficient algorithms for determining μ, except for conservative plants.[3,4] This is directly related to the fact that the quadratic norm yields smooth spherical cost surfaces, whereas the infinity norm yields cubical surfaces that have sharp corners. In practice, σ-analysis is faster and more reliable than Monte Carlo simulation. It is faster than the guaranteed branch-and-bound algorithm of Ref. 5, which would be useful for checking final control designs.

The design method of the paper is based on the SANDY code of Ly[6] and Ly et al.,[7] which designs a single LQG controller, of specified order, for a set of plants with different plant parameters. Here we take this set to be plants having the worst parameter changes for increasing values of σ_i, weighting the more unlikely plants with smaller weights W_i. The SANDY code also only finds a local extremum since it solves a nonconvex problem.

First we review LQG controller design for multiple plants using the SANDY code, with application to the first benchmark problem. Next we describe a modification of LQW σ-analysis to handle LQG problems. Finally we give the new minimax algorithm that combines SANDY and σ-analysis and apply it to the second benchmark problem.

Received Aug. 29, 1991; revision received Dec. 6, 1991; accepted for publication Dec. 18, 1991.

*Ph.D. Candidate, Department of Aeronautics and Astronautics. Senior Member AIAA.

†Pigott Professor of Engineering, Department of Aeronautics and Astronautics. Honorary Fellow AIAA.

Optimal Control Design for Multiple Plants Using SANDY

LQG control gives excellent nominal performance but is generally not robust to parameter variations because it tends to notch lightly damped modes. Optimization of an LQG performance index weighted over several plants produces controllers that are more robust to parameter variations than nominal LQG control. Increasing robustness is especially straightforward when only one parameter varies because it is easy to assess system performance over the entire parameter range.

The computer program SANDY[6,7] uses gradient methods and a nonlinear programming algorithm[8,9] to find the optimal controller for an LQG performance index weighted over N plants:

$$J = \sum_{i=1}^{N} W_i J_{\mathrm{LQG}_i} \tag{1}$$

where W_i is the weight of the ith plant and J_{LQG_i} is the LQG performance index for the ith plant. The sum of the weights is one. SANDY incorporates extensive generality, handling both white and colored noise inputs, reduced-order control, and linear controller constraints. The algorithm is extremely robust; it uses a finite-time performance index, so that stability during all iterations is not necessary for convergence.

Application to Benchmark Problem 1

The two-mass/spring benchmark problem is described in Ref. 10. The system and nomenclature are shown in Fig. 1. The goal of the first benchmark problem is to design a constant-gain linear, noncollocated controller given nominal parameters $m_1 = m_2 = m_{\mathrm{nom}}$ and $k = k_{\mathrm{nom}}$, a measurement of x_2 only, and the following specifications:

1) Provide stability for $0.5 < \bar{k} < 2.0$, $\bar{k} \triangleq k/k_{\mathrm{nom}}$.
2) With nominal parameters, settle x_2 in "about" $15\sqrt{m_{\mathrm{nom}}/k_{\mathrm{nom}}}$ for a magnitude I_w impulsive disturbance.
3) Tolerate "reasonable" measurement noise.
4) Use "reasonable" controller effort, complexity, and bandwidth.

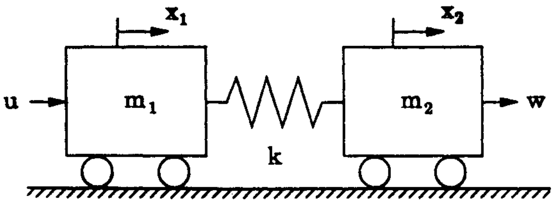

Fig. 1 Nomenclature for two-mass/spring benchmark system.

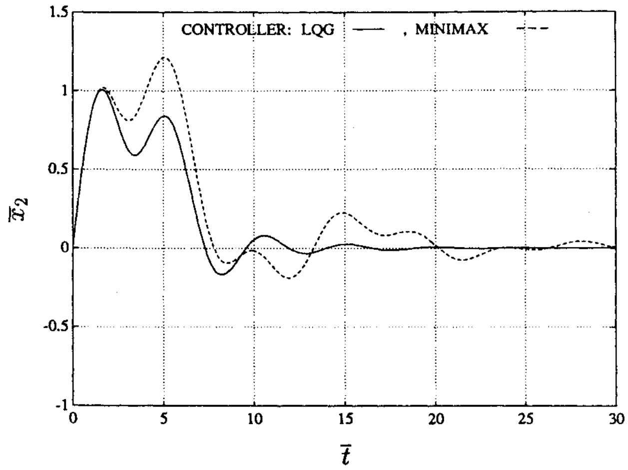

Fig. 2 Benchmark problem 1 response to impulse disturbance I_w with nominal plant parameters.

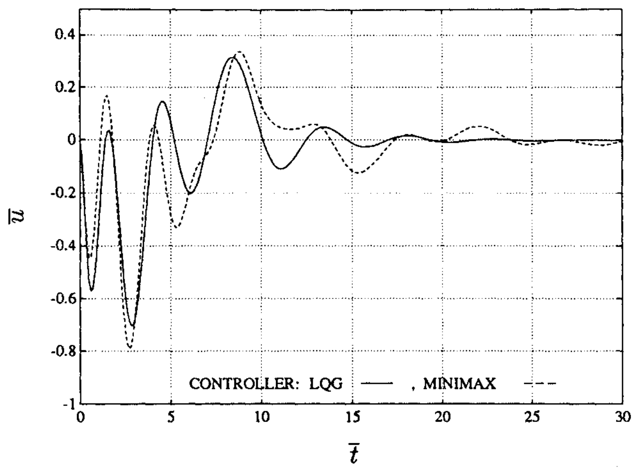

Fig. 3 Benchmark problem 1 control for impulse disturbance I_w with nominal plant parameters.

An overbar denotes normalization by given nominal information, as indicated in specifications 1 and 2. Thus,

$$\bar{t} \triangleq t\sqrt{\frac{k_{\text{nom}}}{m_{\text{nom}}}}, \quad \bar{x}_2 \triangleq x_2\sqrt{k_{\text{nom}} m_{\text{nom}}}\Big/I_w, \quad \bar{u} \triangleq u\sqrt{\frac{m_{\text{nom}}}{k_{\text{nom}}}}\Big/I_w \tag{2}$$

with the same normalization for other quantities with corresponding dimensions.

For this paper, reasonable measurement noise is considered by including it in all performance indices. A conservative ratio of disturbance to measurement noise is

$$Q_n = 2R_n \tag{3}$$

where Q_n and R_n are the spectral densities of the normalized (disturbance, measurement) white noise inputs $\bar{w}$ and $\bar{v}$.

Reasonable controller effort implies that the control magnitudes not exceed one normalized unit and that control energy should be included in all performance indices. This is accomplished by setting

$$Q = \text{diag}(30, 100)R \tag{4}$$

where Q and R are the performance index weights for $[\bar{x}_2, \dot{\bar{x}}_2]$ and $\bar{u}$. Reasonable controller complexity is taken to be full (fourth) order or less, and reasonable bandwidth is taken to be that given by nominal LQG control.

The first benchmark problem was solved using SANDY. The state-space equations of motion for the plant are easily derived and are given in Ref. 10. Since specification 2 states a performance objective for the nominal plant only, we heavily weight the nominal plant, as shown in Table 1.

For the state-space controller form

$$\dot{x}_c = A_c x_c + B_c y_s \tag{5}$$

$$u = C_c x_c \tag{6}$$

where x_c is the controller state and y_s is the measurement including white noise v, SANDY returns:

$$A_c = \begin{bmatrix} -0.7501 & 2.3959 & 0 & 0 \\ -3.0797 & -0.0120 & 0 & 0 \\ 0 & 0 & -2.1027 & 1.1150 \\ 0 & 0 & 6.9133 & -6.4124 \end{bmatrix}$$

$$B_c = \begin{bmatrix} -2.2217 & 0.6943 & 2.4921 & -16.4100 \end{bmatrix}^T$$

$$C_c = \begin{bmatrix} -0.9133 & 1.9696 & -2.4555 & 0.9693 \end{bmatrix} \tag{7}$$

Application of the guaranteed hypercube analysis of Refs. 3 and 4 shows that the controller stabilizes the system for $0.50<\bar{k}<2.00$, meeting specification 1. The nominal LQG controller provides stability for $0.72<\bar{k}<1.57$. Controllers designed for just $\bar{k}=1$ and 0.5 are unstable for $\bar{k}=2$. Similarly, controllers designed for just $\bar{k}=1$ and 2 are unstable for $\bar{k}=0.5$. Therefore, the controller given in Eq. (7) is in effect a minimax controller. It is the optimal controller for the worst parameter changes, which in this case are trivial since there is only one parameter variation.

Figure 2 shows the normalized position of the right mass vs normalized time for a magnitude I_w impulsive disturbance with nominal LQG control and with the minimax controller. The LQG controller meets specification 2. The minimax controller also meets the specification, given reasonable definitions of settling time and "about" in specification 2, but not as well as the nominal LQG controller. However, specifications 1 and 4 mandate performance degradation. Robustness to parameter variations trades against nominal performance.

Figure 3 shows the corresponding control forces. The normalized magnitudes are less than one but, more importantly, the maximum magnitude with the minimax controller is less than 13% higher than that with LQG control. The nominal system bandwidth is 1.51 $\sqrt{k_{\text{nom}}/m_{\text{nom}}}$ with nominal LQG control compared with 1.50 $\sqrt{k_{\text{nom}}/m_{\text{nom}}}$ with the minimax controller.

Worst Parameter Changes

Design of a minimax controller for benchmark problem 1 is straightforward because the worst parameter change is obvious (there is only one parameter variation). Before address-

Table 1 Benchmark plant weightings

Normalized stiffness $\bar{k}$:	0.504	1.000	1.925
Plant weight:	0.001	0.998	0.001

ing minimax control design with two or more parameter variations, it is necessary to solve the problem of finding the direction in the parameter space that maximizes the performance index for a given parameter norm and controller.

Worst parameter changes and control design are treated in Ref. 2. There the worst parameter changes are found for the worst deterministic initial conditions, without process or sensor disturbances. A new parameter maximization problem, which includes white noise disturbance inputs, is

$$J_w = \max_{\Delta p} E\left[\lim_{t_f \to \infty} \frac{1}{2t_f} \int_0^{t_f} (y^T Q y + u^T R u)\, \mathrm{d}t\right] \tag{8}$$

subject to

$$\dot{x} = A(p)x + B(p)u + \Gamma(p)w \tag{9}$$

$$y = C(p)x + D(p)u \tag{10}$$

$$y_s = C_s(p)x + D_s(p)u + \Gamma_s(p)v \tag{11}$$

$$\dot{x}_c = A_c x_c + B_c y_s \tag{12}$$

$$u = C_c x_c + D_c y_s \tag{13}$$

$$E\left[w_a(t) w_a^T(t')\right] = Q_{n_a}\delta(t - t') \tag{14}$$

$$E\left[w_a(t)\right] = 0 \tag{15}$$

$$\left[w_a, x(0)\right] \quad \text{statistically independent} \tag{16}$$

$$w_a \triangleq \begin{bmatrix} w \\ v \end{bmatrix} \tag{17}$$

$$\sigma^2 = \Delta p^T \Sigma^{-2} \Delta p \tag{18}$$

$$\Delta p = p - p_{\text{nom}} \tag{19}$$

where p is a vector of variable parameters. The controller of Eqs. (12) and (13) is *given*. The plant and controller initial conditions are unimportant because their forcing effects are infinitesimal relative to the continual white noise inputs. In the analogous finite-time problem, the initial plant state covariance is important.

Equation (8) is the usual infinite-time LQG performance index, except for the *maximization* with respect to plant parameter variations (hence the subscript w for "worst"). Equations (9) through (16) are standard. Equations (18) and (19) define a hyperellipsoid of norm σ about the nominal point in the parameter space. Thus, the solution to the maximization problem is the point on the hyperellipsoid (direction in the parameter space) at which the performance index is worst for a *given* quadratic norm σ. The relationship of σ to Doyle's μ in Ref. 1 is discussed in Ref. 2. The weighting matrix Σ can have a Gaussian interpretation. Typically it is a diagonal matrix of standard deviations of the parameters about nominal.

The solution to this worst parameter change problem is developed in the Appendix. The resulting necessary conditions for a maximum are

$$0 = A_a X_a + X_a A_a^T + \Gamma_a Q_{n_a} \Gamma_a^T \tag{20}$$

$$0 = A_a^T \Lambda_a + \Lambda_a A_a + Q_a \tag{21}$$

$$\Delta p = \frac{\sigma \Sigma^2 (\partial J/\partial p)^T}{\sqrt{(\partial J/\partial p)\Sigma^2(\partial J/\partial p)^T}} \tag{22}$$

where the gradient of the performance index with respect to the jth parameter is

$$\frac{\partial J}{\partial p_j} = \mathrm{tr}\left[\frac{1}{2} X_a \frac{\partial Q_a}{\partial p_j} + \left(\frac{\partial A_a}{\partial p_j} X_a + \frac{\partial \Gamma_a}{\partial p_j} Q_{n_a} \Gamma_a^T\right)\Lambda_a\right] \tag{23}$$

and the matrices A_a, Γ_a, and Q_a (subscript a denoting augmented plant/controller system) are given information and are defined in the Appendix. The solution is independent of controller order.

Equations (20–23) do not have an apparent closed-form solution, even for simple problems. However, the following homotopic algorithm has worked on all attempts (so far):

1) Select a parameter-change norm σ.
2) Guess Δp of length σ along one of the parameter axes.
3) Solve Eqs. (20), (21), and (23) with the current Δp.
4) Update Δp using Eq. (22).
5) Iterate steps 3 and 4 until Δp converges.
6) Repeat steps 2 through 5 for each positive/negative parameter axis (two cases for each parameter).
7) Record the Δp that produces the worst performance from step 5.

An alternative to testing along the parameter axes (step 2) is to use the gradient at the nominal point or the previous σ: guess $\Delta p = (\partial J/\partial p)\Delta\sigma$.[2] This results in much faster convergence, to the global maximum for many problems. However, for benchmark problem 2 with LQG control, this method fails to find a near-reversal in the worst parameter space direction, from decreasing to increasing spring stiffness as σ is increased past a critical value. Instead, it finds a local maximum associated with decreasing stiffness for all σ. In contrast, an even more reliable but slower method for higher dimension parameter spaces is to guess at the vertices. For example, a two-space has four ($\pm$) axes and four vertices, while a three-space has six axes but eight vertices. Another alternative is random guessing. However, the given method has proven reliable to date.

Application to the Benchmark Problem

The worst parameter algorithm is applied to the benchmark problem for illustration and to discuss reliability in finding the global maximum. To facilitate graphical presentation we allow variations in m_2 and k but assume m_1 is known exactly. The worst parameter changes are found for a range of norms σ with nominal LQG control. This does not constitute a solution to the problem but is only meant to illustrate the worst parameter algorithm. The solution to benchmark problem 2, in which m_1, m_2, and k are uncertain, is presented in the section on minimax control.

We choose different LQG weights for this example than for benchmark problem 1 because the new weights better illustrate the worst parameter algorithm. The new weights are

$$Q = 10R \tag{24}$$

where Q and R are the performance index weights for $\bar{x}_2$ and $\bar{u}$. The velocity $\dot{x}_2$ is not weighted. The white noise disturbance spectral densities are the same as for benchmark problem 1. In particular, these weights yield a less robust nominal LQG controller with a discontinuity in the worst direction in the parameter-change space.

Figure 4 shows the reciprocal of the worst performance index vs the parameter-change norm σ [Eq. (18)] with LQG control. The reciprocal of the performance index is plotted to achieve a reasonable vertical scale. The parameter vector is

$$p = (m_2, k) \tag{25}$$

The weighting matrix in Eq. (18) is chosen to be

$$\Sigma = \mathrm{diag}(m_{\text{nom}}, k_{\text{nom}}) \tag{26}$$

This choice is atypical in that it would usually be a diagonal matrix of standard deviations but this choice gives a more physical interpretation to σ. For example, the parameter margin of 0.18 might correspond to a decrease of 0.18 in one of the normalized parameters (actually both change from nominal).

Figure 5 shows contours of constant cost J over the largest stable circle in the parameter-change space. Contours of con-

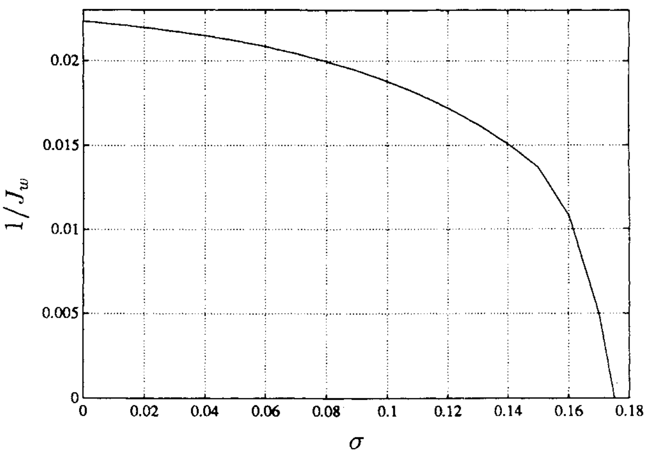

Fig. 4 Benchmark problem worst performance vs parameter-change norm with nominal LQG control.

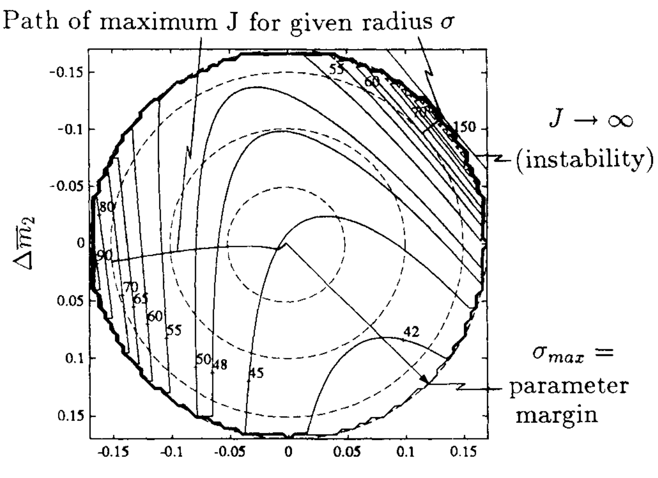

Fig. 5 Benchmark problem cost contours in parameter-change plane with nominal LQG control.

stant parameter-change norm σ appear as circles. The points of highest cost for each norm are shown and produce a path of maxima from the nominal plant at the circle center to the edge where instability first occurs. The path of maxima is discontinuous in general. In this case one discontinuity occurs, between $\sigma = 0.15$ and 0.16. There are two local maxima for each norm, corresponding to two ridges of cost emanating from the nominal point. The discontinuity occurs at the norm for which the height of one ridge becomes higher than the other for a given norm. A corresponding mesh plot showing cost as height above the parameter-change plane is given in Fig. 6.

In general, the worst parameter algorithm should find the global maximum because the two-norm Δp constraint of Eq. (18) yields relatively smooth cost surfaces. An unfound pathological global maximum (a spike of the performance index in the parameter space) is not of great practical importance. The chance of the actual parameters lying at this point would probably be negligible. Also, in many applications, relatively robust, low performance control logic is used as a backup if instability is sensed. The chance of both the nominal and backup controllers being unstable at the same undiscovered point combined with the small chance of the parameters actually lying at that point would inspire confidence in the design.

The benchmark problem solution shown in Fig. 4 was successfully subjected to extensive random checking, which is currently the widely used analysis method in practice. The approximate parameter margin was verified using the hypercube method presented in Refs. 3 and 4. These works show that, with a given controller, conservative systems become unstable at a finite set of closed-loop resonant frequencies (crossings of the $j\omega$ axis). They give the procedure for finding the guaranteed largest parameter-space hypercube, centered at the nominal parameters, in which the closed-loop conservative system is stable. They show that an expanding hypercube first hits the instability region at a corner of the hypercube, corresponding to one of the closed-loop resonant frequencies. For the benchmark problem with LQG control, the normalized distance from the center to the edge of the largest hypercube is 0.125, corresponding to a closed-loop resonant frequency of 1.65 $\sqrt{k_{nom}/m_{nom}}$. Although Eq. (18) defines a hyperellipsoid in general, in this case it is a circle with radius σ. This implies that the largest circle in which the system is stable has a radius between 0.125 and 0.177 (see Fig. 7). The parameter margin of 0.175 indicated in Fig. 4 is therefore consistent with the hypercube analysis.

Minimax Control

Given some desired stable region in the parameter-change space and a corresponding probability distribution of the plant parameters about nominal, the controller should minimize the expected value of the performance index:

$$J_{min} = \min_{A_c, B_c, C_c, D_c} E\left[J(p)\right] \tag{27}$$

where the expected value is taken over all plants in the region (say a hyperellipsoid). As the number of parameters becomes high, just estimating $E[J(p)]$ becomes computationally intensive. This calculation might be impractical for just two or three parameters. Therefore, solving Eq. (27) clearly seems intractable in practice.

A more attainable, but still useful, goal is to find the controller that minimizes an expected value over the *worst plants only*. This corresponds to finding the best controller over the path of maxima shown in Fig. 5 (of course this path is controller dependent). Now the expected value is taken over a line only, regardless of the dimension of the parameter space. For a given number of parameters, the worst path can be found with a high degree of reliability for considerably less effort than that needed to find $E[J(p)]$. If the resulting performance is acceptable over the worst path, it will also be

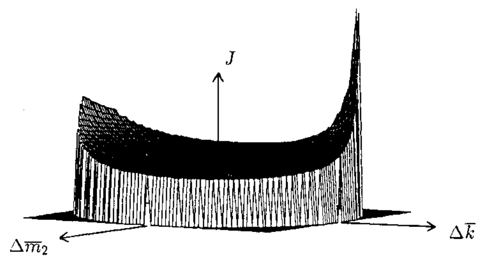

Fig. 6 Benchmark problem cost as height above parameter-change plane with nominal LQG control.

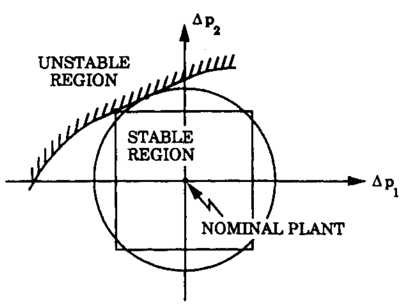

Fig. 7 Comparison of stability square/circle.

acceptable for all plants in the specified region of the parameter-change space. Thus, the goal is to find the controller that *minimizes* the performance index over a set of plants that *maximize* it.

Minimax control combines multiple-plant optimal control design using SANDY with worst parameter changes. Minimax controllers are the optimal controllers for a weighting of worst plant parameter changes. The minimax problem statement is

$$J_s = \min_{A_c, B_c, C_c, D_c} \sum_{i=1}^{N} W_i \max_{\Delta p} J_i \tag{28}$$

where the summation is over N parameter-change norms (worst case plants) including $\sigma = 0$ (the nominal plant). The LQG performance index in Eq. (28) for each plant is

$$J = E\left[\lim_{t_f \to \infty} \frac{1}{2t_f} \int_0^{t_f} (y^T Q y + u^T R u)\, dt\right] \tag{29}$$

Thus, Eq. (28) is just

$$J_s = \min_{A_c, B_c, C_c, D_c} \sum_{i=1}^{N} W_i J_w(\sigma_i) \tag{30}$$

with J_w given in Eq. (8). The constraints on each worst plant are Eqs. (9–11), (18), and (19). Equation (28) is also constrained by Eqs. (12–17). The summation in Eq. (28) could, in principle, be replaced by an integral over σ, from $\sigma = 0$ to some prescribed upper limit, with weight W_i replaced by a probability density for σ. However, the summation is more useful in practice. A similar minimax problem for linear quadratic regulators (full state feedback) has been proposed and solved in Ref. 11.

The solution to the minimax problem depends on an alternate LQG derivation that assumes the controller form a priori. This derivation, which is more straightforward than usual LQG derivations, is given in Ref. 12 for the general time-varying case. The minimax solution is also developed in Ref. 12. Defining

$$X_a \triangleq \begin{bmatrix} X & X_2 \\ X_2^T & X_c \end{bmatrix} \tag{31}$$

$$\Lambda_a \triangleq \begin{bmatrix} \Lambda & \Lambda_2 \\ \Lambda_2^T & \Lambda_c \end{bmatrix} \tag{32}$$

$$Q_{n_a} \triangleq \begin{bmatrix} Q_n & N_n \\ N_n^T & R_n \end{bmatrix} \tag{33}$$

with X_a and Λ_a given by Eqs. (20) and (21) for each *worst* plant, the necessary conditions for a saddlepoint (without feedthrough terms D, D_s, and D_c and with $\Gamma_s = I$) are

$$0 = \sum_{i=1}^{N} W_i (X_{2_i}^T \Lambda_{2_i} + X_{c_i} \Lambda_{c_i}) \tag{34}$$

$$B_c = -\left[\sum_{i=1}^{N} (W_i \Lambda_{c_i})\right]^{-1} \sum_{i=1}^{N} W_i \big[(\Lambda_{2_i}^T X_i + \Lambda_{c_i} X_{2_i}^T) C_{s_i}^T + \Lambda_{2_i}^T \Gamma_i N_n\big] R_n^{-1} \tag{35}$$

$$C_c = -R^{-1} \sum_{i=1}^{N} W_i B_i^T (\Lambda_i X_{2_i} + \Lambda_{2_i} X_{c_i}) \left[\sum_{i=1}^{N} (W_i X_{c_i})\right]^{-1} \tag{36}$$

These conditions are independent of controller order. The necessary conditions with direct feedthrough and $\Gamma_s \neq I$ are more complicated and are given in Ref. 12.

Equations (20), (21), and (34–36) are coupled and a closed-form solution is not apparent. An algorithm that combines SANDY with the worst parameter algorithm and works for any controller order is as follows:

1) With the nominal LQG controller, find the parameter margin σ_{max} using the worst parameter algorithm.

2) Select parameter-change norms $\sigma_1 = 0$, σ_2, $\sigma_3, \ldots$, ($\sigma_i < \sigma_{max}$).

3) Find the worst plant for each σ_i with the current controller using the worst parameter algorithm.

4) Find the optimal controller weighted over the plants from step 2 using SANDY.

5) Iterate steps 2–4 until the necessary conditions are satisfied; if instability results:

   Use current marginally stable worst plant in future weightings.

   Decrease σ_{max} accordingly.

6) Iterate steps 2–5.

   Increase σ_{max} when possible, following step 5.

   Stop when necessary conditions are satisfied and σ_{max} > desired parameter margin.

At certain points in the algorithm, the updated controller can have a lower parameter margin than the previous controller because the closed-loop system becomes unstable in a different direction of the parameter space. For example, for benchmark problem 1, the only variable parameter is the spring stiffness k. The parameter margin with nominal LQG control using the weights in Eqs. (24) and (3) corresponds to a decrease in k. If SANDY is used to find an optimal controller for a weighting of a marginally stable decreased-k plant and the nominal plant, the resulting new closed-loop system will have a larger parameter margin, still corresponding to a decrease in k. However, as this process is continued, after a few iterations the new closed-loop system will have a *smaller* parameter margin corresponding to an *increase* in k. The algorithm will bounce between decreasing k and increasing k forever, without increasing the parameter margin as desired, unless a marginally stable increased-k plant is retained in the weighting as indicated in step 5.

A general MATLAB code was written to implement the minimax algorithm. The code calls SANDY in the form of a MATLAB MEX file,[13] which is compiled, to increase efficiency.

Application to Benchmark Problem 2

Benchmark problem 2 is identical to benchmark problem 1 except for the first specification: instead of providing stability for a specified spring stiffness range, the system must be robust over an unspecified range of three parameters: the two masses and the spring stiffness (see Fig. 1). The minimax algorithm was used to solve benchmark problem 2 by heavily weighting the nominal plant and lightly weighting a worst plant within 1% of the parameter margin and the plants kept from step 5 of the algorithm. Initially, the nominal plant was weighted 0.99 and the worst plant near the current stability boundary weighted 0.01. The nominal plant weighting was lowered to 0.98 once retained marginally stable plants appeared in step 5. The extra 0.01 of weighting was then divided equally between the retained plants from step 5. The actual choice of weightings is not particularly important because, in general, severe weighting changes are required to produce significant performance and robustness differences for the resulting minimax controller.

The LQG weights used for benchmark problem 2 are those given in the worst parameter change example. As indicated there, these weights provide a more challenging and interesting test of the worst-parameter algorithm, and therefore the minimax algorithm, than the weights used for benchmark problem 1. The benchmark problem 1 weights provide a significantly shorter settling time for x_2.

A minimax controller that solves benchmark problem 2 is

$$A_c = \begin{bmatrix} -0.3508 & 1.9311 & 0 & 0 \\ -2.6114 & -0.3506 & 0 & 0 \\ 0 & 0 & -1.6418 & 1.0438 \\ 0 & 0 & 0.3718 & -1.9417 \end{bmatrix}$$

$$B_c = \begin{bmatrix} -1.6924 & 0.2150 & 1.8365 & -3.9936 \end{bmatrix}^T$$

$$C_c = \begin{bmatrix} -0.4654 & 1.2739 & -2.0440 & 0.8660 \end{bmatrix} \qquad (37)$$

The reciprocal cost vs parameter-change norm σ with this minimax controller is the dashed curve in Fig. 8. The solid curve is the relation with nominal LQG control. The minimax controller doubles the parameter margin from nominal LQG and is relatively insensitive to parameter variations over the entire LQG stability range. Moreover, the nominal performance (inverse cost) is decreased only 8.9%. The nominal performance must decrease somewhat because it trades against robustness and parameter margin. The minimax parameter margin of 0.31 indicated in Fig. 8 is consistent with the guaranteed hypercube analysis of Refs. 3 and 4 as described in the "Worst Parameter Changes" section.

An alternate minimax controller that triples the parameter margin is

$$A_c = \begin{bmatrix} -0.3508 & 1.9311 & 0 & 0 \\ -3.2929 & -0.3245 & 0 & 0 \\ 0 & 0 & -1.6418 & 1.0438 \\ 0 & 0 & 1.8567 & -2.6038 \end{bmatrix}$$

$$B_c = \begin{bmatrix} -1.4652 & 0.0299 & 1.5533 & -4.2474 \end{bmatrix}^T$$

$$C_c = \begin{bmatrix} -0.4654 & 1.2739 & -2.0440 & 0.8660 \end{bmatrix} \qquad (38)$$

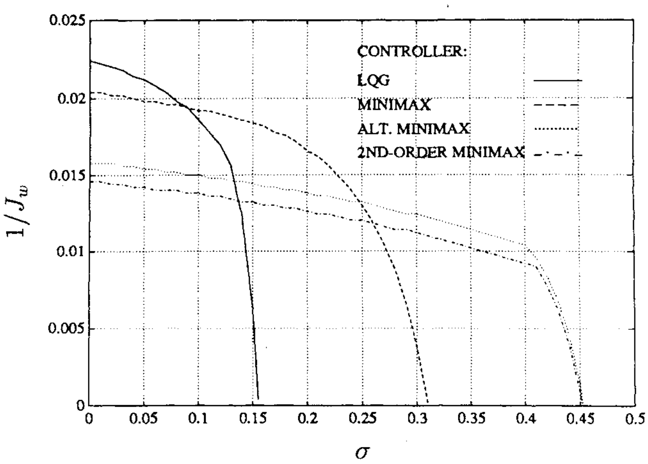

Fig. 8 Benchmark problem 2 worst performance vs parameter-change norm: controller comparison.

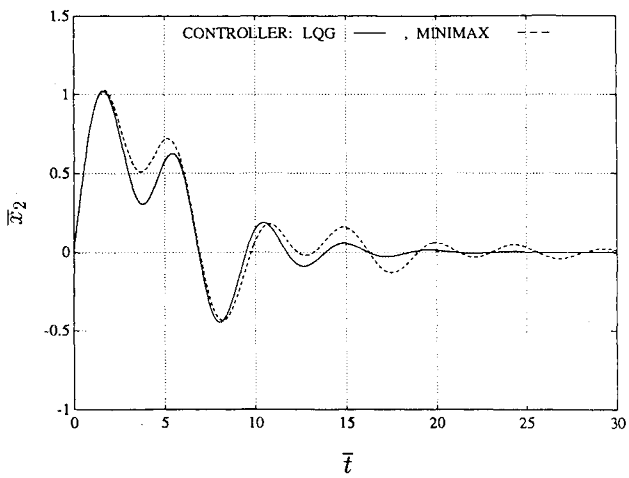

Fig. 9 Benchmark problem 2 response to impulse disturbance I_w with nominal plant parameters.

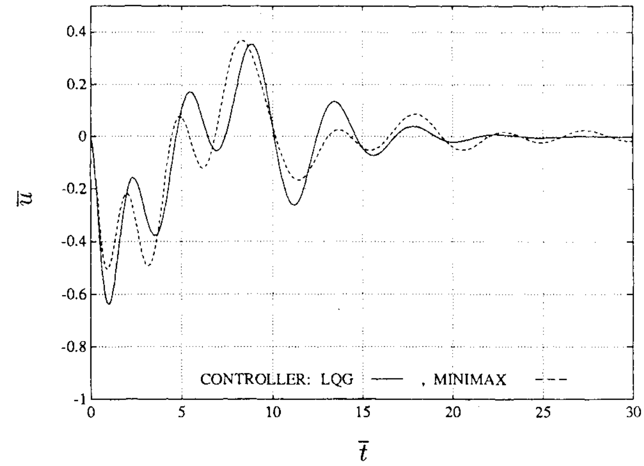

Fig. 10 Benchmark problem 2 control for impulse disturbance I_w with nominal plant parameters.

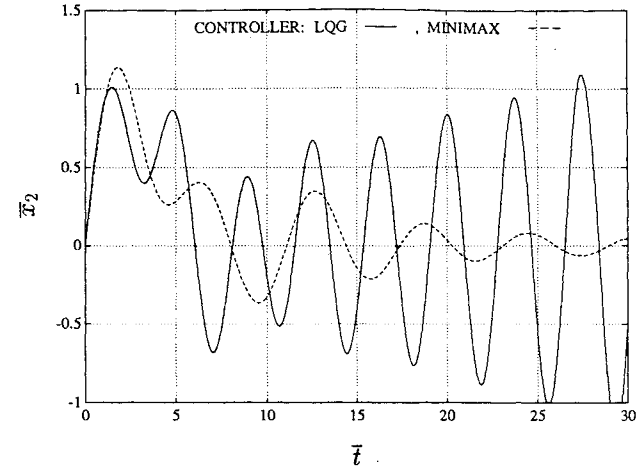

Fig. 11 Benchmark problem 2 response to impulse disturbance I_w with worst plant parameter changes of norm 0.2.

The reciprocal cost vs parameter-change norm σ with this minimax controller is the dotted curve in Fig. 8. Although this controller is also relatively insensitive to parameter variations, the nominal performance is decreased 29% from LQG. The parameter margin of 0.45 indicated in Fig. 8 is consistent with the hypercube analysis of Refs. 3 and 4.

A reduced-order minimax controller that triples the parameter margin is

$$A_c = \begin{bmatrix} -0.7500 & 0.5809 \\ -0.2314 & -1.0594 \end{bmatrix}$$

$$B_c = \begin{bmatrix} -0.0816 & 2.0039 \end{bmatrix}^T$$

$$C_c = \begin{bmatrix} 0.5000 & -0.4648 \end{bmatrix} \qquad (39)$$

The reciprocal cost vs parameter-change norm σ with this minimax controller is the dash-dotted curve in Fig. 8. This controller exhibits the same insensitivity to parameter variations as the corresponding full-order minimax controller of Eq. (38) with only a 7.6% drop in nominal performance. Again, the parameter margin of 0.45 indicated in Fig. 8 is consistent with the hypercube analysis of Refs. 3 and 4.

Figure 9 shows the normalized position of the right mass vs normalized time for a magnitude I_w impulsive disturbance with nominal LQG control and with the minimax controller of Eq. (37), with nominal plant parameters. The minimax controller performs almost as well as the LQG controller and meets the settling-time requirement, depending on the settling-

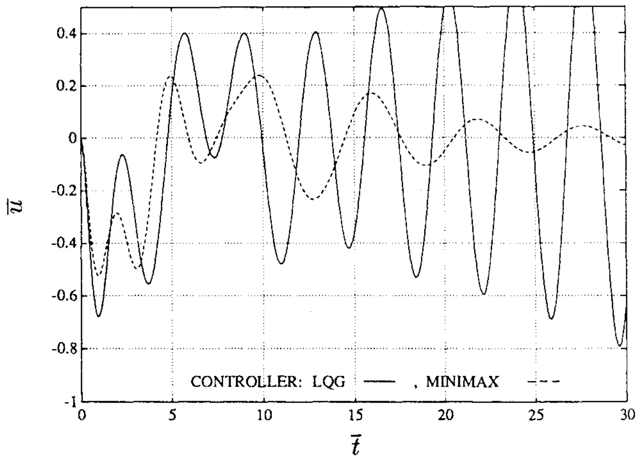

Fig. 12 Benchmark problem 2 control for impulse disturbance I_w with worst plant parameter changes of norm 0.2.

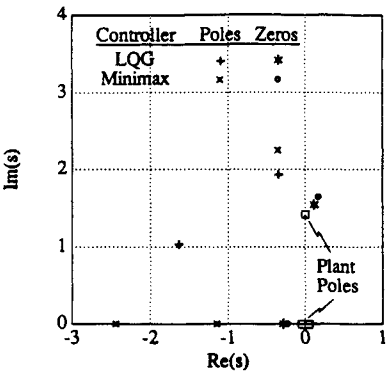

Fig. 13 Benchmark problem 2 controller pole/zero locations.

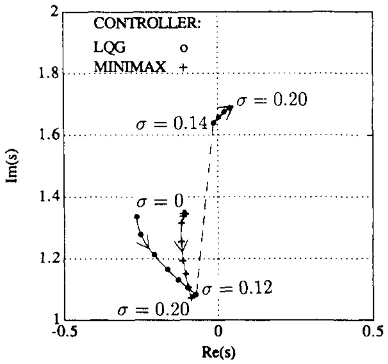

Fig. 14 Benchmark problem 2 closed-loop plant oscillatory mode vs parameter-change norm σ.

time definition. Figure 10 shows the corresponding control forces. The normalized magnitudes are less than one, and the maximum magnitude with the minimax controller is the same as with LQG control.

The minimax controller outperforms the LQG controller for worst plants with parameter-change norm greater than 0.09. Figures 11 and 12 show the normalized right-mass position and normalized control force histories, respectively, for the worst $\sigma = 0.2$ plants. The worst parameter changes are shown in Table 2.

The worst parameter algorithm has found a 146-deg change in the worst parameter space direction with LQG control from $\sigma = 0.1$. The system is unstable with LQG control, whereas the minimax controller is still performing well relative to its nominal performance. The LQG parameter changes for this case were obtained by extrapolating the worst parameter change direction from the marginally stable norm $\sigma = 0.15$. The LQG-controlled system is actually unstable over a broad range of the parameter space for $\sigma = 0.2$.

Table 2 Worst parameter changes for $\sigma = 0.2$

Controller	$\Delta\bar{p}_w \triangleq (\Delta\bar{m}_1, \Delta\bar{m}_2, \Delta\bar{k})$
LQG	$(-0.0955, -0.0996, 0.1448)$
Minimax	$(0.0278, -0.0077, -0.1979)$

Figure 13 shows the pole/zero locations of the LQG and minimax controllers. The nominal plant poles are also shown (there are no plant zeros). Both the LQG and minimax controllers combine an LQG lead (a zero near $s = 0$ and two poles further left) to compensate for the plant rigid-body mode, with a nonminimum phase notch, to compensate for the plant vibration mode. Relative to the LQG notch, the minimax notch is detuned slightly with respect to the plant vibration mode, decreasing the sensitivity to changes in this mode. The change in the zeros from LQG to minimax is less than that of the poles, but a unit of zero change has more effect than a unit of pole change because the zeros are much closer to the plant poles. The most significant change is the movement of two complex LQG controller poles to the real axis with minimax control.

The pole/zero placement shows that the minimax controller is not a new type of controller. However, the minimax control design algorithm has placed the lead and notch very intelligently to increase robustness. Although an experienced designer might duplicate this performance for single-input/single-output (SISO) plants, the minimax algorithm handles multiple-input/multiple-output (MIMO) plants equally well, where even an experienced designer could not.

The nominal system bandwidth is $1.47\sqrt{k_{\text{nom}}/m_{\text{nom}}}$ with LQG control vs $1.48\sqrt{k_{\text{nom}}/m_{\text{nom}}}$ with the minimax controller. Figure 13 indicates that the Bode magnitude roll-off characteristics are similar with each controller, with the same ultimate roll-off rate. Thus, the minimax controller has not decreased robustness to unmodeled dynamics while increasing robustness to parameter variations.

Figure 14 shows the closed-loop system root locus for the plant oscillation mode vs the parameter-change norm σ, for σ from 0 to 0.2 in 0.02 increments. This is the slowest closed-loop mode (the LQG loci move much more than the minimax loci for the faster modes). There is a discontinuity in the locus with LQG control due to the 146-deg change in the worst parameter space direction after $\sigma = 0.12$. The superior nominal performance of the LQG controller is evident in the faster placement of the nominal closed-loop plant oscillation mode. However, the LQG oscillation mode dives towards instability as σ varies. In contrast, the minimax controller trades the 8.9% drop in nominal performance for an oscillation mode that rides parallel to the imaginary axis, maintaining stability and performance.

Minimax control designs for a six-parameter helicopter and a four-parameter space telescope[12] perform similarly relative to nominal LQG control.

Conclusions

Minimax controllers are optimal controllers for a set of worst plant parameter changes. Necessary conditions for a minimax have been derived using a new LQG derivation. A minimax controller synthesis code combines an optimal multi-plant control design code, SANDY, with a worst parameter algorithm that uses a quadratic norm on parameter changes. The worst parameter algorithm finds the global maximum of the performance index in the parameter space with a high degree of confidence, withstanding extensive random checking, but statements of parameter margin are not guaranteed. However, for a two-mass/spring benchmark problem, all parameter margin results are consistent with a guaranteed-stability hypercube method for conservative systems.

Minimax controllers, including a reduced-order controller, were designed to solve two benchmark problems for the two-mass/spring system. The first benchmark problem has one parameter variation whereas the second problem has three parameter variations. The results validate the trade of parameter margin and parameter insensitivity against nominal performance. Relative to nominal LQG control, minimax control doubles the parameter margin, is insensitive to parameter variations, and shows insignificant control magnitude and bandwidth differences, while decreasing nominal performance less than 10%. Minimax control does not decrease robustness to unmodeled dynamics in increasing robustness to parameter variations. Thus, minimax control provides near-optimal nominal performance with significant robustness and parameter margin improvements.

Appendix: Worst Parameter Problem Solution

To solve the worst parameter changes problem given by Eqs. (8–19), first augment x with x_c:

$$x_a \triangleq \begin{bmatrix} x \\ x_c \end{bmatrix} \tag{A1}$$

It follows that

$$\dot{x}_a = A_a x_a + \Gamma_a w_a \tag{A2}$$

where

$$A_a \triangleq \begin{bmatrix} A + BD_c\bar{D}C_s & B\bar{D}_cC_c \\ B_c\bar{D}C_s & A_c + B_c\bar{D}_sC_c \end{bmatrix} \tag{A3}$$

$$\Gamma_a \triangleq \begin{bmatrix} \Gamma & BD_c\bar{D}\Gamma_s \\ 0 & B_c\bar{D}\Gamma_s \end{bmatrix} \tag{A4}$$

$$\bar{D} \triangleq (I - D_sD_c)^{-1} \tag{A5}$$

$$\bar{D}_c \triangleq I + D_c\bar{D}D_s \tag{A6}$$

and $\bar{D}$ is assumed to exist.

Using the identity $x_a^TQ_ax_a \equiv \mathrm{tr}(x_ax_a^TQ_a)$, where tr is the trace operator, and realizing the integrand in Eq. (8) will reach steady state, after some algebra,

$$J = \tfrac{1}{2}\,\mathrm{tr}(X_aQ_a) \tag{A7}$$

where

$$X_a \triangleq E(x_ax_a^T) \tag{A8}$$

and the blocks of the symmetric matrix Q_a are

$$Q_{a_{11}} = \bar{C}^TQ\bar{C} + \bar{C}_s^TR\bar{C}_s \tag{A9}$$

$$Q_{a_{12}} = (\bar{C}^TQD + \bar{C}_s^TR)\bar{D}_cC_c \tag{A10}$$

$$Q_{a_{22}} = C_c^T\bar{D}_c^T\bar{R}\bar{D}_cC_c \tag{A11}$$

where

$$\bar{C} \triangleq C + DD_c\bar{D}C_s \tag{A12}$$

$$\bar{C}_s \triangleq D_c\bar{D}C_s \tag{A13}$$

$$\bar{R} \triangleq D^TQD + R \tag{A14}$$

and the well-posedness assumption that $D_c\bar{D}\Gamma_s = 0$ has been made. This assumption insures that $E(vv^T)$, which is infinite for Gaussian white noise, does not appear in Eq. (A7).

The constraint of Eq. (18) is adjoined to the performance index with a Lagrange multiplier ρ. Another constraint is the well-known (e.g., Ref. 14) equation, Eq. (20), for the steady-state augmented-state covariance X_a. This n_a-dimension Lyapunov equation can be written as an n_a^2-length vector equation since it is linear. The vector equation can then be adjoined to Eq. (A7) with a vector of Lagrange multipliers. However, it is equivalent and simpler to use a matrix of Lagrange multipliers Λ_a and the trace operator:

$$\bar{J} \triangleq \tfrac{1}{2}\Big\{\mathrm{tr}(X_aQ_a) + \mathrm{tr}\big[(A_aX_a + X_aA_a^T + \Gamma_aQ_{n_a}\Gamma_a^T)\Lambda_a\big] - \rho(\Delta p^T\Sigma^{-2}\Delta p - \sigma^2)\Big\} \tag{A15}$$

Now take the differential of $\bar{J}$ by taking the differential of X_a and the quantities that depend on the plant parameters, using the trace identities $\mathrm{tr}(AB) = \mathrm{tr}(BA)$, $\mathrm{tr}(A^T) = \mathrm{tr}(A)$, and $\mathrm{tr}(A + B) = \mathrm{tr}(A) + \mathrm{tr}(B)$ repeatedly:

$$d\bar{J} = \tfrac{1}{2}\,\mathrm{tr}\big[\mathrm{d}X_a(Q_a + \Lambda_aA_a + A_a^T\Lambda_a)\big] + \mathrm{tr}\left\{\left[\frac{1}{2}X_a\frac{\partial Q_a}{\partial p} + \left(\frac{\partial A_a}{\partial p}X_a + \frac{\partial \Gamma_a}{\partial p}Q_{n_a}\Gamma_a^T\right)\Lambda_a\right]\mathrm{d}p\right\} \tag{A16}$$

Setting the coefficients of $\mathrm{d}X_a$ and $\mathrm{d}p$ to zero, respectively, yields the remaining necessary conditions for a maximum, Eq. (21) and

$$\Delta p = \frac{1}{\rho}\Sigma^2\left(\frac{\partial J}{\partial p}\right)^T \tag{A17}$$

with $(\partial J/\partial p_j)$ defined in Eq. (23). Substituting Eq. (A17) into Eq. (18) yields Eq. (22).

The finite-time derivation follows immediately with the inclusion of initial conditions and nonzero time derivatives.

Acknowledgment

This research was supported by the United States Air Force Laboratory Graduate Fellowship Program.

References

[1]Doyle, J. C., "Analysis of Feedback Systems with Structured Uncertainty," *IEEE Proceedings,* Vol. 129, Pt. D, No. 6, 1982, pp. 242–250.

[2]El Ghaoui, L., Carrier, A., and Bryson, A. E., "Linear-Quadratic Minimax Controllers," *Journal of Guidance, Control, and Dynamics,* Vol. 15, No. 4, 1992, pp. 953–961.

[3]El Ghaoui, L., and Bryson, A. E., "Worst Parameter Changes for Stabilized Conservative Systems," *Proceedings of the AIAA Guidance, Navigation, and Control Conference* (New Orleans), AIAA, Washington, DC, Aug. 1991.

[4]El Ghaoui, L., "Robustness of Linear Systems to Parameter Variations," Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, March 1990.

[5]Balakrishnan, V., Boyd, S., and Balemi, S., "Branch and Bound Algorithm for Computing the Minimum Stability Degree of Parameter-Dependent Linear Systems," *International Journal of Robust and Nonlinear Control* (submitted for publication).

[6]Ly, U.-L., "A Design Algorithm for Robust Low-Order Controllers," Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Stanford Univ., (SUDDAR 536), Stanford, CA, Nov. 1982.

[7]Ly, U.-L., Cannon, R. H., and Bryson, A. E., Jr., "Design of Low-Order Compensators Using Parameter Optimization," *Automatica,* Vol. 21, No. 3, 1985, pp. 315–318.

[8]Gill, P. E., and Murray, W., "Quasi-Newton Methods for Unconstrained Optimzation," *Journal of Mathematics and Its Applications,* Vol. 9, 1972, pp. 91–108.

[9]Gill, P. E., Murray, W., and Wright, M., *Practical Optimization,* Academic Press, New York, 1981.

[10]Bernstein, D. S., and Wie, B., "Benchmark Problems for Robust Control Design," 1991 American Control Conference (proposal for an invited session), Oct. 1991.

[11]Lau, M. K., Boyd, S., Kosut, R. L., and Franklin, G. F., "Robust Control Design for Ellipsoidal Plant Set," *Proceedings of the 30th Conference on Decision and Control* (Brighton, England, UK), Dec. 1991.

[12]Mills, R. A., "Parameter-Robust Controller and Estimator Design Using Minimax Methods," Ph.D. Thesis, Dept. of Aeronautics and Astronautics, Stanford Univ., Stanford, CA, April 1992.

[13]*PRO-MATLAB User's Guide,* MathWorks, Inc., South Natick, MA, Jan. 1990.

[14]Bryson, A. E., and Ho, Y.-C., *Applied Optimal Control,* Hemisphere, Washington, DC, 1975.