Efficient Treatment of Moderate Amplitude Constraints for Helicopter Handling Qualities Design Optimization

Vineet Sahasrabudhe* and Roberto Celi†
University of Maryland, College Park, Maryland 20742

This paper describes a new technique for the calculation of gradients of constraints associated with the moderate amplitude criteria of the ADS-33 helicopter handling qualities specifications. The gradients are calculated using low-order linear approximations to the full nonlinear model of the helicopter. The low-order models approximate the gradients well and reduce the additional cost of calculating the gradient by a factor of about 50. Most of the reduction in the objective function obtainable using the exact gradients are retained, with no additional infeasible intermediate designs. The accuracy of linear Taylor-series expansions of the constraint in terms of the design variables is found to depend on the size of the changes of each design variable. The accuracy is not improved by using intermediate design variables, such as reciprocals and cubes of the design variables, but it improves if the bandwidth or derivative ratios such as $L_{\theta_{1c}}/L_p$ are used as intermediate variables in the expansions.

Nomenclature

$F(X)$	= objective function value
$G_{L\text{base}}(s)$, $G_{L\text{pert}}(s)$	= baseline and perturbed low-order transfer functions
$G(s)$, $G_{\text{pert}}(s)$	= baseline and perturbed transfer functions
$g_j(X)$	= jth constraint
k_β, k_ζ	= flap and lag spring stiffnesses for main rotor blades
k_ϕ, k_p	= roll attitude and rate feedback gains
L_p, $L_{\theta_{1c}}$	= roll stability and control derivatives
p_{pk}	= peak roll rate
p, q, r	= aircraft angular velocities in roll, pitch, and yaw
R	= flap-lag blade elastic coupling parameter
u	= input displacement vector
u, v, w	= aircraft velocity components along the body axes
W_1, W_2	= frequency response weighting functions
X	= vector of design variables
$(X_i)_U$, $(X_i)_L$	= upper and lower bound on ith component of design vector
x	= vector of states
$x^{(j)}$	= flight control system state
$\beta_0, \beta_{1c}, \beta_{1s}, \beta_2$	= rotor flap degrees of freedom in nonrotating coordinate system
$\Delta\phi_{\text{min}}$, ϕ_{pk}	= minimum and peak roll attitudes
$\zeta_0, \zeta_{1c}, \zeta_{1s}, \zeta_2$	= rotor lag degrees of freedom in nonrotating coordinate system
$\lambda_j(\cdot)$	= jth eigenvalue
$\lambda_0, \lambda_{1c}, \lambda_{1s}$	= inflow degrees of freedom
ϕ, θ, ψ	= aircraft attitudes in roll, pitch, and yaw

Introduction

A RECENT trend in helicopter design is the emphasis on high maneuverability, with handling qualities tailored to specific tasks, to increase mission effectiveness and decrease pilot workload. The steps required to achieve this may lead to stiff rotors and high-gain flight control systems, which strongly couple rotor/fuselage and flight control system dynamics. It is reasonable to assume that an integrated, rather than independent, design of rotor and flight control systems will provide better performance and reduce the risk of closed-loop instabilities. On the other hand, an integrated approach leads to complex design problems, with large numbers of design variables and constraints. One tool that may help deal with such complexity is the use of formal numerical optimization techniques. Reference 1 describes a study for the simultaneous optimization of rotor and flight control system subject to aeroelastic stability and handling quality constraints. A simple coupled rotor–fuselage helicopter model with a model-following flight control system was considered in the study.

This paper addresses a problem identified in Ref. 1, namely, the large computational effort required to solve integrated rotor/control systems optimization problems of realistic complexity. Formulations with 10 design variables and 30–40 constraints can require CPU times of the order of 70–80 h on typical workstations. A major contributor to the problem is the handling quality specifications that call for the calculation of time histories of the aircraft response to pilot inputs. This includes the criteria for moderate and large attitude change maneuvers in the ADS-33 Handling Qualities specifications.[2] Most algorithms for constrained optimization require the calculation of the gradients of the constraints with respect to the vector of design variables. The gradients are typically calculated using finite difference approximations, which require that the time histories of the aircraft response be calculated as many times as there are design variables every time a gradient is needed. This can translate into hundreds of such calculations for problems of even modest complexity.

Several methods have been developed to improve the computational efficiency of optimizations involving constraints based on time histories. One class of methods relies on analytic and semianalytic expressions for the gradients. There is a large body of work on the response sensitivity of linear systems, first addressed by Tomović.[3] The methods for linear models take advantage of the linearity of the governing equations for both response and sensitivity. As a result, it is possible to generate sensitivities with respect to any number of system parameters for the cost of evaluating just one additional system model.[4] Adelman and Haftka[5] give an overview of several methods of transient response sensitivity analysis. When there are many more active constraints than design variables, the

direct method is used. In this case, the sensitivity equations are solved together with the dynamic response equations; computational savings are obtained because of the similar form of the two sets of Eq. (6). When there are a few constraints and many design variables, an adjoint method is preferred.[7] Finally, when the number of response quantities used to calculate the constraints is smaller than the number of constraints or design variables, the Green's function method is the most efficient.[8] Several of these methods could be applied in the present study. However, they tend to require the integration of systems of differential equations that may be large or expensive to set up. Therefore, it may be useful to explore alternative methods better tailored to the specific features of the problem.

A different approach is to replace the original dynamic response constraint with an equivalent constraint. Feng et al.[9] did this by integrating the constraint over the time interval. This constraint is potentially nondifferentiable and can cause numerical difficulties. In Hsieh and Arora,[10] the constraint is replaced by a pointwise response constraint enforced at critical time points, which are those times at which the constraint is locally close to being infeasible. Grandhi et al.[11] present various techniques to reduce the computational effort required to identify the critical points. These techniques are not immediately useful in the present study because they address a somewhat different problem. In fact, the moderate amplitude constraints require that the time histories have a certain shape, rather than enforcing rigid boundaries on the amplitude, such as limits on its peaks. Also, the shape is defined by parameters that are available after the first peak and the first valley of the time-history plot. Therefore, any algorithm used to identify the position of subsequent peaks or valleys, such as those of Ref. 11, does not significantly help reduce the computational effort.

Another class of methods deals with the high computational cost by creating a sequence of approximate problems, the solutions of which converge to the optimum of the original problem. This can be accomplished by carrying out Taylor-series expansions of objectives and constraints about the current design point[13] in terms of the design variables themselves or of intermediate variables. Applications of this and other approximation concepts to structural optimization are reviewed in Ref. 12. Linear or, more rarely, quadratic Taylor-series expansions in terms of the design variables are implemented as part of general purpose optimization software.[14,15] Taylor-series expansions in terms of intermediate variables[6] have been used infrequently outside structural optimization, because these design variables are discipline and problem dependent. In particular, little or no research has been devoted to the identification of suitable intermediate variables for many helicopter-related problems.

This paper has the following objectives:

1) To describe a new technique for an efficient calculation of the gradients of constraints associated with moderate attitude change specifications in ADS-33. This technique is based on calculating approximate values of the gradients from appropriate low-order, linear approximations to the complete nonlinear dynamics of the helicopter.

2) To present results of an integrated rotor-flight control system optimization with aeroelastic and handling qualities constraints, in which the exact moderate amplitude constraints in pitch and roll are replaced by the approximations mentioned earlier.

3) To explore the replacement of moderate amplitude constraints with polynomial approximations based on Taylor-series expansions and to determine what expansions are most likely to be accurate over a wide region of the design space.

Formulation of the Optimization Problem

The helicopter model consists of a rigid fuselage and a main rotor composed of four rigid blades with offset hinges and root springs. Quasisteady aerodynamics with a dynamic inflow model is used. The coupled rotor–fuselage–inflow equations are written in first-order form as

$$\dot{x} = f(x, u, t) \tag{1}$$

where x and u contain the pitch settings for the main and tail rotors. A small perturbation model is obtained by linearizing Eq. (1) about a trimmed equilibrium position. This results in

$$\Delta\dot{x} = A\Delta x + B\Delta u \tag{2}$$

where Δ denotes a small perturbation from trim. This is the model used for frequency response calculations. The flight control system architecture is similar to that of the UH-60 ADOCS (Ref. 16), based on a model-following concept.

The optimization is formulated to seek a vector X of design variables that minimizes an objective function $F(X)$, subject to behavior constraints $g_j(X) \leq 0$, $j = 1, \ldots, m$, and side constraints $(X_i)_L \leq X_i \leq (X_i)_U$, $i = 1, \ldots, N$, where m is the number of constraints, N is the number of design variables, and $(X_i)_L$ and $(X_i)_U$ are the limits on the design variable X_i. Ten design variables are used in the study: the root spring stiffnesses of the rotor, the flap-lag elastic coupling parameter R (Ref. 17), selected feedback gains, and selected poles in the command model that describes the desired dynamics of the helicopter.

Aeroelastic and handling qualities constraints are imposed on the design. Aeroelastic constraints enforce rotor aeroelastic stability by requiring that the real parts of λ_j of the linearized state matrix A corresponding to the rotor modes be negative. The handling qualities constraints enforce compliance with a representative subset of the ADS-33 specifications[2] for hover. Two types of objective function are used. When the optimizer cannot obtain a feasible design, the objective function is formulated as the sum of the values of the violated constraints. When the design is feasible, the objective function chosen is a weighted sum of the swashplate displacements and rates for two predefined moderate amplitude maneuvers; this weighted sum is representative of the control effort. Additional details concerning the analysis model and the design optimization problem can be found in Ref. 1.

A limitation of this study is that actuator saturation is not included in the aircraft model. Saturation plays a role in the ability to satisfy the moderate amplitude response specifications and, therefore, it should be considered for a complete treatment of the problem. On the other hand, Ref. 1 showed that saturation was unlikely to occur for the configurations analyzed in that study. Because the same configurations are analyzed in the present study, the limitations of the model should not affect significantly the results presented in this paper.
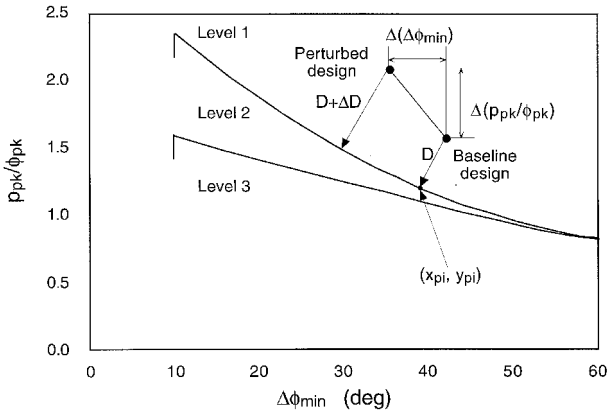
Treatment of Moderate Amplitude Constraints

Definitions

The moderate amplitude attitude change specifications (or quickness specifications) need to be transformed into inequality constraints before they can be used in the optimization. Only the treatment of one specification for roll is described here. This is specification 3.3.3, for the target acquisition and tracking mission task element in hover and low-speed flight.[2] The constraints associated with the pitch specifications are treated similarly.

The first step is to define the distance $\mathcal{D}$ between the point representative of the helicopter on the specification chart and the curve that limits the level 1 region (Fig. 1). This distance is given by

$$\mathcal{D} = \sqrt{(x_{pi} - \Delta\phi_{min})^2 + [y_{pi} - (p_{pk}/\Delta\phi_{pk})]^2} \tag{3}$$

where $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ are the x and y coordinates of the representative point on the chart, and x_{pi} and y_{pi} are the coor-

Fig. 1 Definition of some quantities used in the formulation of the constraints and their gradients.

dinates of the intersection of the perpendicular from that point to the level 1 curve. The quantities $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ are determined from the helicopter response to a step-like lateral cyclic input, ϕ_{pk} and p_{pk} are the maximum roll angle and roll rate following the pilot input, respectively, and $\Delta\phi_{min}$ is the roll angle of the first valley of the response (or the final value if the response is not oscillatory).[2] The specification is transformed into a constraint $g(X)$ by requiring that $g(X) = \alpha\mathcal{D} \leq 0$, where $\alpha = -1$ if the point is in the level 1 region, and $\alpha = +1$ if it is not. The gradient of $g(X)$ is given by

$$\nabla g(X) = \left\{\frac{\partial g}{\partial X_i}\right\} = \left\{-\frac{\partial\mathcal{D}}{\partial X_i}\right\}$$

$$= -\left\{\frac{\partial\mathcal{D}}{\partial(p_{pk}/\Delta\phi_{pk})}\frac{\partial(p_{pk}/\Delta\phi_{pk})}{\partial X_i} + \frac{\partial\mathcal{D}}{\partial\Delta\phi_{min}}\frac{\partial\Delta\phi_{min}}{\partial X_i}\right\} \quad (4)$$

where X_i is the ith design variable, and the size of the vector is equal to the total number of design variables N. Several quantities that appear in Eq. (4) are shown graphically in Fig. 1. The terms $\partial\mathcal{D}/\partial(p_{pk}/\Delta\phi_{pk})$ and $\partial\mathcal{D}/\partial\Delta\phi_{min}$ are simple geometric quantities that can be calculated numerically very quickly. The terms $\partial(p_{pk}/\Delta\phi_{pk})/\partial X_i$ and $\partial\Delta\phi_{min}/\partial X_i$ are the changes in the specification parameters, because of changes in the design variables, for a fixed pilot input. These terms may be expensive to calculate because they depend on the changes of the time histories of the aircraft response to pilot inputs. Therefore, an objective of the present study is to explore strategies to calculate these sensitivities more efficiently.

The assumption will be made that the sensitivities of the state matrix A and of the control matrix B in Eq. (2) are known. It is reasonable to assume that in an optimization with handling qualities constraints these sensitivities would have to be calculated anyway to evaluate the effects on the small-amplitude change specifications. In that case, baseline and perturbed frequency responses would be available for use with the moderate amplitude constraints considered in this study at no additional cost. Efficient methods to calculate the sensitivities of the linearized system for helicopter problems are presented in Ref. 18.

Approximations Using Low-Order Linear Models

Low-Order Transfer Function Fits

The quickness specifications are closely related to the bandwidth of the helicopter, which is the frequency at which the phase margin is 45 deg or the gain margin is 6 dB (Ref. 2). For example, for a simple first-order model of roll rate dynamics, the parameter $p_{pk}/\Delta\phi$ is equal to the bandwidth. Also, the quickness requirements blend with the small-amplitude (bandwidth) requirements as the attitude changes approach the lower values of the definition of moderate. Whether or not a heli-

copter satisfies the quickness specifications is heavily influenced by its frequency response characteristics in a band that roughly extends from 1 to 10 rad/s. Therefore, the idea is to calculate approximate values of $\partial(p_{pk}/\Delta\phi_{pk})/\partial X_i$ and $\partial\Delta\phi_{min}/\partial X_i$ in Eq. (4) from a low-order linear model that fits the frequency response of the helicopter over that limited frequency band. The procedure is composed of the following steps:

1) Calculate the frequency response of the baseline configuration and fit a low-order model over a suitable frequency band, e.g., from 0.5–1 to 10–15 rad/s.

2) Integrate the low-order model to find the response to a prescribed pilot input, and determine $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$. These will be the baseline values for the purpose of calculating the gradients.

3) Perturb X_i, calculate the frequency response of the perturbed configuration, and again fit a low-order model over the chosen frequency band.

4) Integrate the perturbed low-order model to find the response to the same pilot input, and determine $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$. These will be the perturbed values.

5) Use baseline and perturbed values to calculate the sensitivities $\partial(p_{pk}/\Delta\phi_{pk})/\partial X_i$ and $\partial\Delta\phi_{min}/\partial X_i$ using finite difference approximations.

The choice of the transfer function types for this study is based on a survey of approximate transfer functions described in Ref. 19, which cover several aircraft sizes and rotor configurations. The roll transfer function can then be written in the form

$$G_L(s) = \frac{\phi(s)}{\theta_{1c}(s)} = \frac{k(s + z)}{(s + p)(s^2 + 2\zeta\omega_n s + \omega_n^2)} \quad (5)$$

or, in shorthand notation:

$$G_L(s) = \frac{\phi(s)}{\theta_{1c}(s)} = \frac{k(z)}{(p)[\zeta; \omega_n]} \quad (6)$$

Here, k, p, z, ζ, and ω_n are variables that are adjusted to fit the frequency response to that of the full model. Two simpler versions of Eq. (5) are also used in this study, namely, a transfer function with three poles and no zeros, and one with two poles and one zero.

Let $G(j\omega)$ denote the baseline frequency response ϕ/θ_{1c}, obtained using the full linearized model of the helicopter. Let $G_{pert}(j\omega)$ be the frequency response of the perturbed design, also calculated using the complete linearized model, and obtained by perturbing one of the design variables. Then, the low-order model frequency response $G_L(j\omega)$ is fitted to $G(j\omega)$ by adjusting the values of the free parameters to minimize the error ε, defined by

$$\varepsilon = \sum_k W_1(\omega_k)\{\text{Re}[G(\omega_k)] - \text{Re}[G_L(\omega_k)]\}^2$$
$$+ \sum_k W_2(\omega_k)\{\text{Im}[G(\omega_k)] - \text{Im}[G_L(\omega_k)]\}^2 \quad (7)$$

where the summations extend over a suitable number of frequency points $j\omega_k$. The weight functions $W_1(\omega_k)$ and $W_2(\omega_k)$ can be used to focus on desired frequency bands. Let $G_{Lbase}(s)$ be the result of the fit to the baseline transfer function. The process is repeated a second time by fitting the low-order transfer function to $G_{pert}(j\omega)$. This results in the approximate $G_{Lpert}(s)$. The time histories required for the determination of the baseline and the perturbed values of $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ are then obtained by integrating the state-space realizations corresponding to $G_{Lbase}(s)$ and $G_{Lpert}(s)$, respectively. The gradients in Eq. (4) are calculated using finite differences.

Low-Order State-Space Models

Low-order models can also be extracted from the full-order linearized model of Eq. (2) using condensation. If the state

vector Δx is partitioned into vectors of states Δx_R to be retained and Δx_C to be condensed out, and the system matrices are partitioned accordingly, one has

$$\begin{Bmatrix} \Delta \dot{x}_R \\ \Delta \dot{x}_C \end{Bmatrix} = \begin{bmatrix} A_{RR} & A_{RC} \\ A_{CR} & A_{CC} \end{bmatrix} \begin{Bmatrix} \Delta x_R \\ \Delta x_C \end{Bmatrix} + \begin{Bmatrix} B_R \\ B_C \end{Bmatrix} \Delta u \qquad (8)$$

Condensation is based on setting $\Delta \dot{x}_C = 0$ and solving for Δx_C. This results in

$$\Delta \dot{x}_R = (A_{RR} - A_{RC} A_{CC}^{-1} A_{CR}) \Delta x_R + (B_R - A_{RC} A_{CC}^{-1} B_C) \Delta u$$

$$= A_1 \Delta x_R + B_1 \Delta u \qquad (9)$$

This method is more convenient to implement than the transfer function fit because it only requires simple matrix manipulations. Additionally, the resulting models are multi-input multi-output and retain cross-coupling information, whereas the low-order fits are single-input single-output only. On the other hand, it is not obvious which states should be retained to obtain good approximations over the frequencies of interest.

The state vector for the full linear model in Eq. (2) consists of 34 states, that is

$$\Delta x = [u \ v \ w \ p \ q \ r \ \phi \ \theta \ \psi \ \beta_0 \ \beta_{1c} \ \beta_{1s} \ \beta_2 \ \dot{\beta}_0 \ \dot{\beta}_{1c} \ \dot{\beta}_{1s} \ \dot{\beta}_2 \ \zeta_0 \ \zeta_{1c}$$

$$\zeta_{1s} \ \zeta_2 \ \dot{\zeta}_0 \ \dot{\zeta}_{1c} \ \dot{\zeta}_{1s} \ \dot{\zeta}_2 \ \lambda_0 \ \lambda_{1c} \ \lambda_{1s} \ x_c^{(1)} \ x_c^{(2)} \ x_c^{(3)} \ x_c^{(4)} \ x_c^{(5)} \ x_c^{(6)}]^T \qquad (10)$$

In this study we consider three low-order linear models obtained through condensation, with 4, 10, and 14 states, respectively. The vector Δx_{R14} of states for the 14-state model is

$$\Delta x_{R14} = [p \ q \ \phi \ \theta \ x_c^{(1)} \ x_c^{(2)} \ x_c^{(3)} \ x_c^{(4)} \ x_c^{(5)} \ x_c^{(6)} \ \dot{\beta}_{1c} \ \dot{\beta}_{1s} \ \beta_{1c} \ \beta_{1s}]^T \qquad (11)$$

The state vector Δx_{R10} for the 10-state model is the same as Δx_{R14}, minus the flap states. The state vector Δx_{R4} for the four-state model is the same as Δx_{R14} minus the flap states and the six flight control system states $x^{(j)}$. These low-order models are used in the same general way as those obtained from frequency response fits. Thus, the corresponding equations of motion are integrated to evaluate the moderate amplitude constraint for both the baseline and the perturbed designs, and the gradients are calculated using finite difference approximations.

The accuracy of the approximate gradients will depend, for both types of low-order models, primarily on three parameters. These are 1) the frequency band over which the low-order model fits the full-order model, 2) the number of poles and zeros (for the low-order fits) or the number and type of states (for the state-space models), and 3) the size of the perturbations used in the numerical calculation of the gradients. Two other parameters may play a role when low-order fits are used, namely, the weighting functions W_1 and W_2 in Eq. (7), and the convergence criterion used to stop the error minimization in Eq. (7). The techniques presented in this section are compatible with any flight dynamic simulation and do not require changes in the computer programs implementing the simulation. Because the quickness specification value tends to depend on the frequency response in the 0.5–1 to 10 rad/s range, the rest of this study concentrates on this range. More precisely, the fit is performed over the interval from 0.3 to 20 rad/s, with the interval from 0.3 to 10 rad/s emphasized using the weighting functions W_1 and W_2. These values have been found to produce the most accurate and repeatable results.

Taylor-Series Expansions

The optimization problem is solved using a feasible sequential quadratic programming (FSQP) algorithm.[14] As part of the solution process used by FSQP, objective function and constraints are expanded in linear or quadratic Taylor-series about the current design. The expansions are updated at each iteration. It is sometimes convenient to perform the Taylor-series

expansions not directly in terms of design variables but rather in terms of intermediate variables that are functions of the design variables. The goal is to obtain higher-quality expansions that are accurate over a larger portion of the design space and, therefore, require less frequent updates (each update of a linear expansion requires a gradient calculation). It would then be possible to relax the limits to the maximum changes of the design variable at each iteration, required to prevent the design from entering portions of the design space where the Taylor-series expansions are too inaccurate.

The FSQP implementation of Ref. 14 is limited to expansions in terms of the design variables, unless the Taylor-series expansions are supplied by the user. A limited amount of experimentation was performed to determine whether the use of intermediate variables could be beneficial. Recall that if X^0 is the current design vector, the linear Taylor-series expansion of a constraint $g(X)$ in terms of the design variables is given by

$$g(X) = g(X^0) + \sum_{i=1}^{N} \frac{\partial g}{\partial X_i} \bigg|_{X_i = X_i^0} (X_i - X_i^0) \qquad (12)$$

If we define a vector Y of intermediate variables generically as $Y_i = h(X_i)$, then the Taylor-series expansion of Eq. (12) can be rewritten in terms of Y as

$$g(Y) = g(Y^0) + \sum_{i=1}^{N} \frac{\partial g}{\partial Y_i} \bigg|_{Y_i = Y_i^0} (Y_i - Y_i^0)$$

$$= g(Y^0) + \sum_{i=1}^{N} \left[\frac{\partial g}{\partial X_i} \frac{\partial X_i}{\partial Y_i} \right] \bigg|_{X_i = X_i^0} [h(X_i) - h(X_i^0)] \qquad (13)$$

Note that $g(Y^0)$ is the value of the constraint at the current design, therefore, $g(Y^0) = g(X^0)$. The intermediate design variables considered in this study are reciprocal design variables, $Y_i = 1/X_i$, and cubic design variables, $Y_i = X_i^3$. Two additional types of Taylor-series expansions have been considered in this study. They can both be written in the general form

$$g(l) = g(l^0) + \frac{dg}{dl} \bigg|_{l=l^0} (l - l^0) \qquad (14)$$

where $l = l(X)$ is an intermediate design variable, and $l^0 = l(X^0)$ is the value of l at the current design. The types of l considered in the study for the roll constraint are

$$l_i(X) = \frac{L_{\theta_{1c}}}{L_p} \qquad l_2(X) = \omega_{BW} \qquad (15)$$

where ω_{BW} is the bandwidth.[2] The stability and control derivatives $L_{\theta_{1c}}$ and L_p are obtained from a six-degree-of-freedom model of the helicopter, derived as shown by Eqs. (8) and (9), by retaining rigid body velocities, rates, and Euler angles. The intermediate variable l_1 has been selected because $-l_1 \theta_{1c}$ is the steady-state value of p following a step input of lateral cyclic of magnitude θ_{1c} in hover, using a simple one-degree-of-freedom linearized roll model. The variable l_2 has been selected because $\omega_{BW} = p_{pk}/\Delta \phi_{min}$ for the same first-order model. In other words, both choices of l correspond to parameters that play an important role in the response of the helicopter. The derivative in Eq. (14) is given by

$$\frac{dg}{dl} \bigg|_{l=l^0} = \sum_{i=1}^{N} \frac{\partial g}{\partial X_i} \frac{\partial X_i}{\partial l} \bigg|_{X=X^0} \qquad (16)$$

The first term in the summation is the ith component of $\nabla g(X)$. The second is the derivative of X_i with respect to l_1 or l_2, and it can be calculated using finite difference approximations. If the sensitivity of Eq. (2) has already been calculated,

then that derivative is already available and can be reused in Eq. (16). Once $g(l^0)$ and $dg/dl|_{l=l^0}$ are available, the constraint for any design X can be calculated directly from Eq. (14) if the relationships $\omega_{BW}(X)$ or $L_{\theta_{1c}}/L_p(X)$ are known. In this study, such relationships are assumed to be available from the consideration of bandwidth constraints. If they are not, then the cost of calculating $l_1 = \omega_{BW}$ or $l_2 = L_{\theta_{1c}}/L_p$ for a given design X must be taken into account. In general, any of these Taylor-series expansions can be used without changes in the computer programs that implement the flight simulation. Finally, when comparing Eqs. (12–14), the gradient $\nabla g(X)$ is calculated numerically from the full nonlinear model of the helicopter; the low-order linear approximations previously described are not used.

Results

The results presented in this section refer to a soft-in-plane hingeless rotor helicopter. The solidity of the main rotor is $\sigma = 0.07$ and its angular velocity is 424 rpm. The weight coefficient is $C_w = 0.005$. The rotor has four blades and the blade airfoil has a lift curve slope of $a = 6.0$ and a profile drag coefficient of $c_{d_0} = 0.01$. The rotor is attached to a fuselage with a gross weight of 4855 lb. The c.g. of the fuselage is $0.2R$ below the the hub.

Approximate Gradients

Figure 2 compares exact and approximate values of three components of the gradient of the roll quickness constraint. The exact values are obtained using the nonlinear model of the helicopter, the approximate values using the linear approximation of Eq. (5), that is a transfer function with three poles and one zero. The finite difference step is of 0.1 and 0.5% for Figs. 2a and 2b, respectively. The components correspond to the following three design variables: 1) $X_1 = k_p = k_\phi$, the roll

attitude and rate feedback gains, linked to be identical; 2) $X_5 = a_1$, one of the poles of the command block in the feedforward path; and 3) $X_8 = k_\zeta$, the stiffness of the lag spring of the rotor blades. Five values are shown for each variable, one for each of five designs corresponding to the first five iterations of an actual optimization. Figure 2 is drawn in such a way that if the exact and the approximate values were identical they would appear on the diagonal of the figure. Also shown are sectors corresponding to relative differences of ± 20 and $\pm 50\%$. The gradient components corresponding to different design variables can have different units and, therefore their respective magnitudes are not directly comparable; this should be kept in mind when examining figures like Fig. 2.

Figure 2 shows that all of the components of the gradient have the correct sign. The best agreement with the exact values from the full nonlinear model is for those with respect to X_1. The worst agreement is for the derivatives with respect to X_8, which are very small; this, together with actual values of the lag frequency that are typically around 0.7–0.8, indicates that the gradient of the constraint would not be very sensitive to changes in X_8. From the results of Fig. 2, and others not presented in this paper, it appears that a step size of 0.1% is the best for the finite difference calculation of the gradients.

Figure 3 compares the roll quickness specification values for perturbed designs calculated using the full nonlinear model and

Table 1   Computer times for gradient calculation

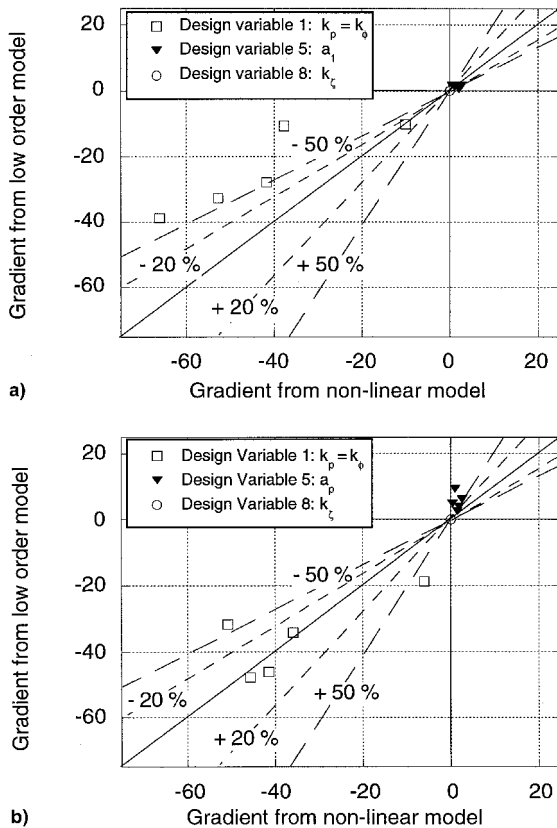| Type of model	CPU time, s	Relative CPU cost, full model = 100
Full nonlinear	308	100.0
Two-pole/one-zero	7	2.3
Four-state	5	1.6
Ten-state	8	2.6
Fourteen-state	12	3.9

Fig. 2 Derivatives of roll moderate amplitude constraint with respect to three design variables for five different designs; comparison between values from full nonlinear model and from three-pole/one-zero low-order model. Perturbation a) 0.1 and b) 0.5%.
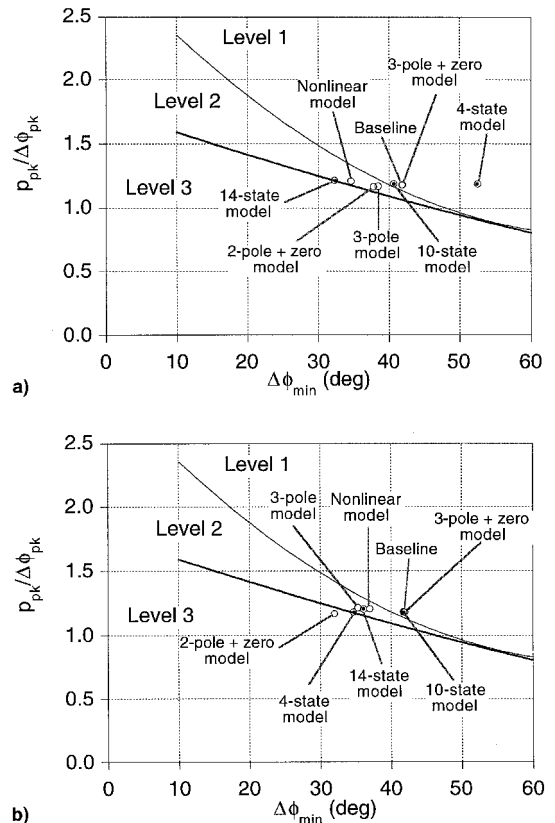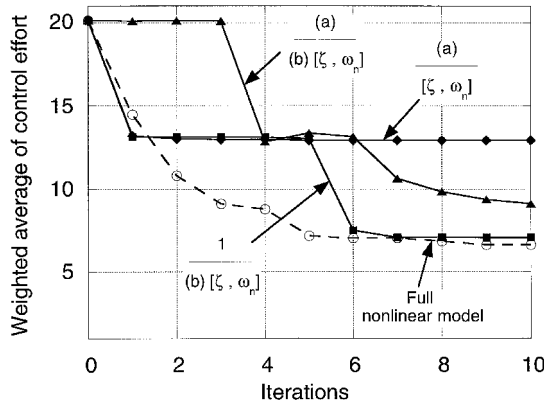

Fig. 3 Roll moderate amplitude specification values for baseline and perturbed designs: a) perturbation of $X_1 = k_p$, and b) $X_5 = a_1$.

Table 2 Design variables: initial and final values

Design variable	Initial design	Nonlinear final design	Linear models — final designs		
			Three poles	Two poles/ one zero	Three poles/ one zero
$k_\beta \times 1.0e\text{-}04$	5.390 (1.12)	5.372 (1.119)	5.533 (1.14)	5.505 (1.13)	5.505 (1.13)
$k_\zeta \times 1.0e\text{-}05$	1.147 (0.70)	1.345 (0.72)	1.249 (0.71)	1.2499 (0.715)	1.2499 (0.715)
R	0.900	0.898	0.921	0.820	0.820
$k_p,\ k_\phi$	−0.0288	−0.0398	−0.02835	−0.2825	−0.2825
k_q	0.5332	0.3546	0.3632	0.4252	0.4252
k_θ	0.0692	0.0619	0.0659	0.1589	0.1589
$k_r,\ k_\psi$	−0.0194	−0.0376	−0.0188	−0.0182	−0.0182
$a_1,\ a_2$	0.8827	0.6125	0.8307	0.8412	0.8412
$a_3,\ a_4$	0.7982	0.8435	0.7913	0.7925	0.7924
Objective	20.102	6.633	7.040	12.959	9.082

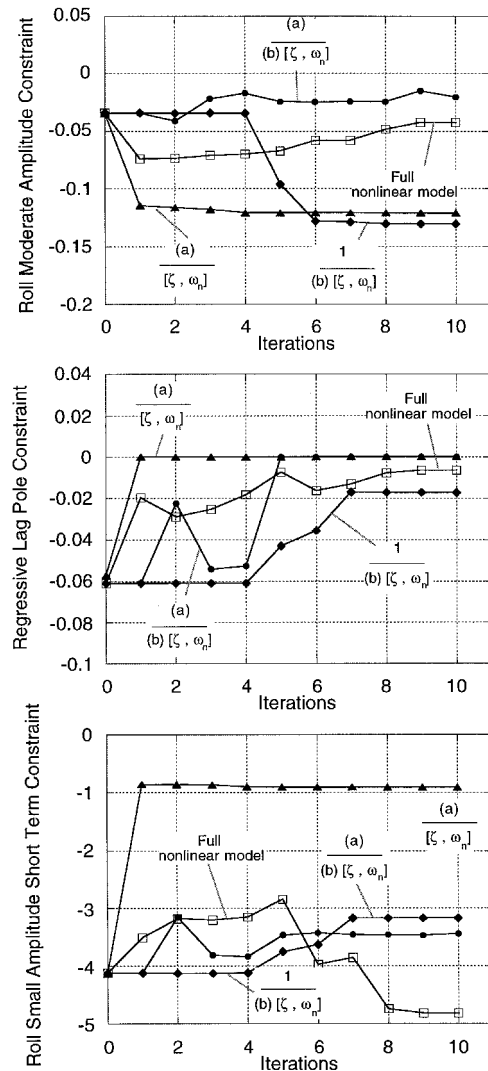[a] Flap and lag frequencies shown in parentheses = in./rev.

Fig. 4 Iteration history of objective function.

the various low-order models. Figure 3a refers to perturbations of X_1, and Fig. 3b refers to perturbation of X_5. The perturbed designs are obtained starting from the baseline design, using the full nonlinear model; next, the relevant components of the gradients are calculated using the full nonlinear model and various low-order approximations. These components are multiplied by a preassigned change of the design variable to obtain the total increments in $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ because of a change in X_1 (upper plot) or X_5 (lower plot). Finally, these increments are added to the baseline values of $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ to obtain the perturbed points shown in the plots. Figure 3a indicates that in most cases the low-order approximations give qualitatively correct predictions, and in some cases the quantitative agreement is also good. The low-order models are used to predict the perturbations only. In fact, using them to also calculate the baseline values of $\Delta\phi_{min}$ and $p_{pk}/\Delta\phi_{pk}$ sometimes gives inaccurate results.

Table 1 compares the average CPU times required to calculate one component of the gradient. These times should be added to that of the baseline constraint evaluation. The additional work necessary to obtain the gradients using the low-order approximations is clearly a minuscule fraction of that required using the full nonlinear model.

Optimization with Approximate Gradients

This section presents the results of an optimization for the case of hover, formulated with control effort as the objective function. The constraint gradients, except for the roll and pitch quickness constraints, are obtained using the full nonlinear model of the helicopter. The gradients of the roll and pitch constraints are calculated using the nonlinear model and the three low-order transfer function fits. When the nonlinear model is used, the results are identical to those of Ref. 1. Figure 4 shows the iteration history of the objective function. All of the points in the plot are from full nonlinear analyses. The four formulations converge to different optima. This is not surprising because the approximate constraints are obtained by

Fig. 5 Comparison of selected constraint iteration histories for optimization using full nonlinear gradients and approximate gradients.

truncating the dynamics of the helicopter and not by a local approximation of the mathematical behavior of the constraint. Therefore, the approximation will not improve as the optimum is reached. The best results are obtained with a transfer function with three poles, but substantial reductions are also obtained using the other low-order fits. The intermediate designs are feasible. Table 2 shows initial and final values of the design variables using approximate and full gradients. The three-pole approximation achieves most of the reduction in objective

function that could be obtained from the full nonlinear model. In practice, any of the low-order approximations could be used at the beginning of the optimization before switching to the full nonlinear model at the end to achieve the full improvement.

Figure 5 shows the iteration histories of some constraints. The roll moderate amplitude constraint is never active or violated. The same is true for the pitch constraint, not shown in the figure. Both are calculated using the full nonlinear model, but their gradients are calculated using the low-order approximations (except for the full nonlinear model curve in the figure). The approximations do not take the optimizer into the infeasible region. The regressive lag stability constraint becomes active when the two-pole/one-zero and the three-pole/one-zero approximations are used, and nearly active when the full model is used. Both this constraint and its gradients are always computed using the full nonlinear model. The results for the approximate models are different because the FSQP optimizer selects different designs, since the shape of the entire feasible region is changed by the approximate constraints. Other optimization algorithms were not explored in this study. However, the behavior would likely be similar if penalty function methods were used, because of their tendency to funnel down the middle of the feasible region. On the other hand, if feasible direction-type methods were used, the iteration histories of objective function, constraints, and design variables

would likely be identical in the exact and the approximate cases, unless the moderate amplitude roll (or pitch) constraint became active. In fact, these methods reach the optimum by moving along the constraints, and only the roll (or pitch) constraint would be changed by the approximations. The roll small amplitude constraint also remains satisfied for all models. The use of the one-zero/two-pole approximation moves the design much closer to the level 1 boundaries than the other approximations; all of them move the design closer than when the full model is used. Figure 6 shows the iteration histories of selected design variables. The optimizer places the desired roll pole a_1 in the feedforward block to a noticeably higher value than when the full model is used. The pitch axis pole shows a much smaller change. The parameter R converges to values between 0.82–0.92.

Taylor-Series Expansion of Roll Moderate Amplitude Constraint

Figure 7 compares the roll constraints calculated from the full nonlinear model and from the linear Taylor-series expansion. Only one design variable at a time is perturbed, i.e., the summation in Eq. (12) contains only one term. For Fig. 7a the variable is k_β, whereas for Fig. 7b it is k_ζ. In each case, the expansions are carried out around four designs, chosen among those that occurred during the optimization runs described in the previous section. Each expansion is evaluated with perturbations $X_i - X_i^0$ of 0.1, 0.5, 1, and 5% of the current value X_i^0 of the design variable (all indicated as small changes), of 10, 20, and 40% (intermediate changes) and 80% (large). Figure 7a shows that linear expansions predict the correct sign of the constraint for all values of k_β. The accuracy is very good for small and intermediate perturbations and reasonable for large perturbations. Hence, relatively large move limits can be allowed for k_β in the roll constraint during the optimization. Figure 7b refers to the design variable k_ζ and shows that for large moves away from the current design the sign of the linear Taylor-series approximation to the constraint can be wrong. This implies that a violation would be predicted when the constraint is satisfied, or vice versa. For changes of small and intermediate size the sign is correct, but the accuracy deteri-
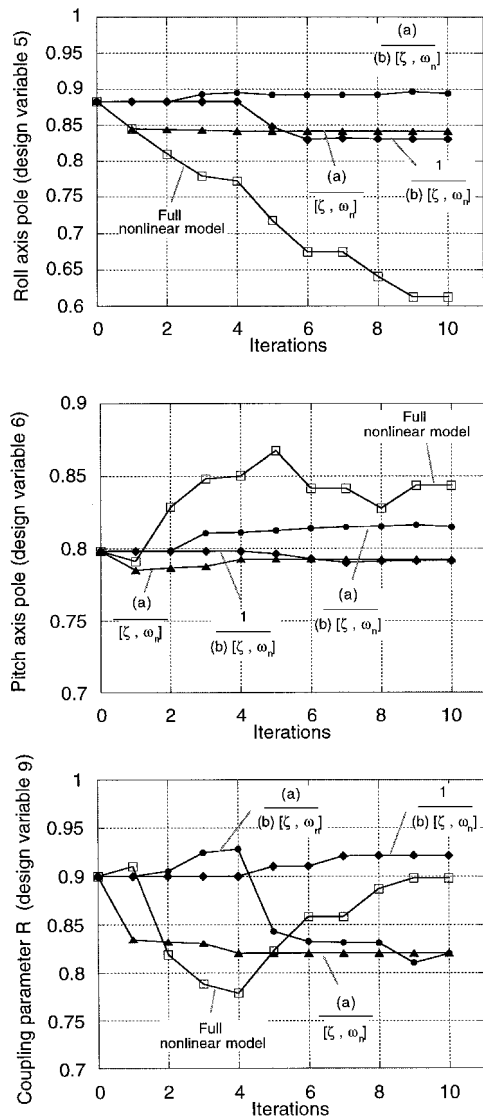

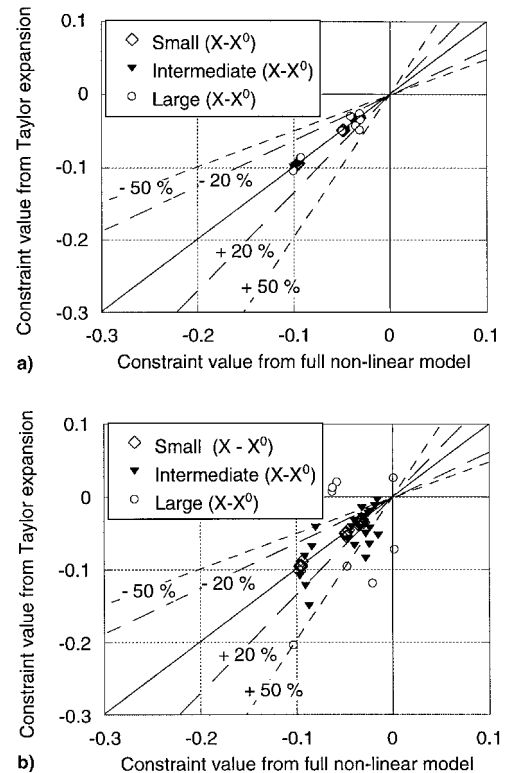Fig. 6   Iteration histories of selected design variables.

Fig. 7   Taylor-series expansions of roll moderate amplitude constraint: a) k_β and b) k_ζ.

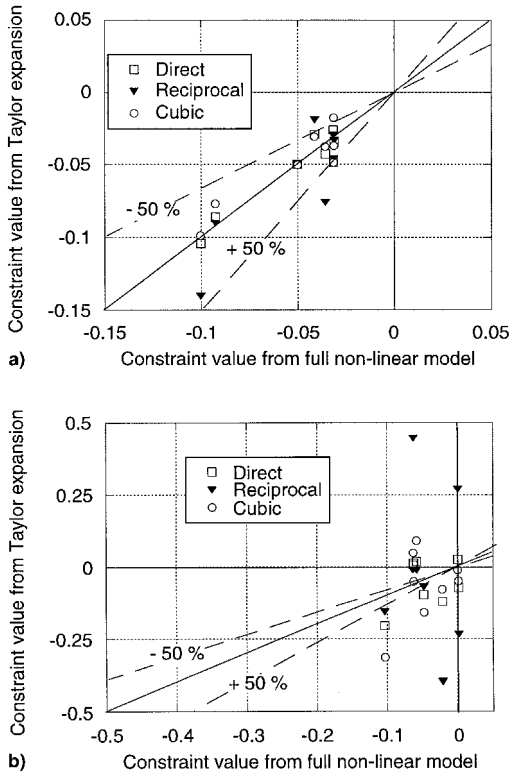Fig. 8 Taylor-series expansion of roll moderate amplitude constraint in terms of direct and intermediate variables: a) k_β and b) k_ζ.
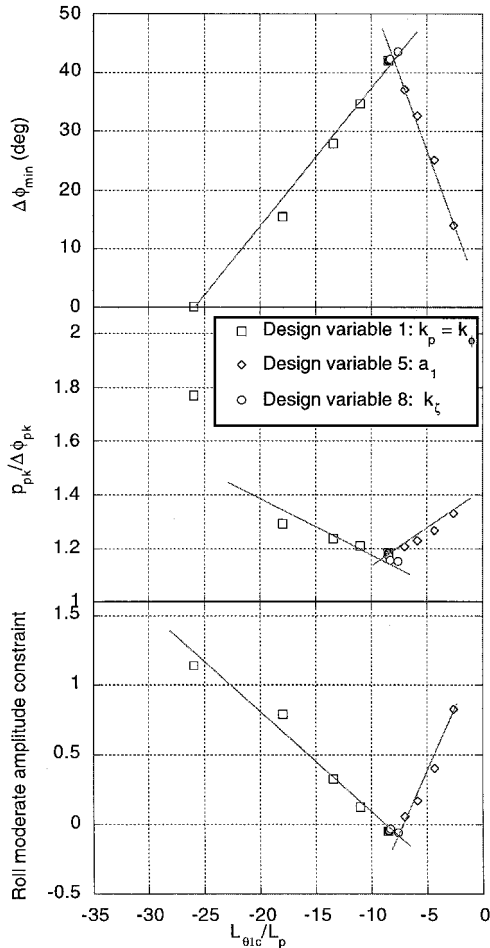
orates as the size increases. The expansions with respect to k_β and k_ζ are representative of the best and the worst performance of the linear Taylor-series approximation to the roll quickness constraint in this problem. On the basis of these results, the recommended maximum move limit is about 40% of the value of each design variable.

Figure 8 again compares the roll quickness constraint predicted by a Taylor-series approximation and by the full nonlinear model. As in Fig. 7, k_β and k_ζ are perturbed. Data for three types of expansions are presented, namely, expansions in terms of direct design variables (same as Fig. 7), of reciprocal, and of cubes of design variables. The perturbations are all 80% of the current value X_i^0, corresponding to the large perturbations of Fig. 7. Using these two types of intermediate variables does not extend the range of validity of the Taylor-series expansions. The situation is different when $l_1 = L_{\theta_{1c}}/L_p$ and $l_2 = \omega_{BW}$ are used. Figure 9 shows some results for the former case. The three plots of Fig. 9 have been constructed by taking a representative design and perturbing three design variables by amounts corresponding to intermediate and large perturbations. Then, $\Delta\phi_{min}$, $p_{pk}/\Delta\phi_{pk}$, $L_{\theta_{1c}}/L_p$, and the roll constraint corresponding to each of the perturbed designs have been calculated and the representative points plotted in the figure. As each design variable is perturbed, most of the points tend to fall on a straight line if $\Delta\phi_{min}$, $p_{pk}/\Delta\phi_{pk}$, and the constraint are plotted as a function of $L_{\theta_{1c}}/L_p$. The points corresponding to changes in $X_8 = k_\zeta$ are collapsed into almost one point in all the plots, but they fall on a straight line when the x axis scale is expanded. The straight lines on the plots have been drawn manually to show the trends. Figure 10 has been obtained exactly as Fig. 9, except that ω_{BW} has replaced $L_{\theta_{1c}}/L_p$. The same trends as in Fig. 9 are evident in Fig. 10. Together, they suggest that

Fig. 9   Roll moderate amplitude constraint for different designs.
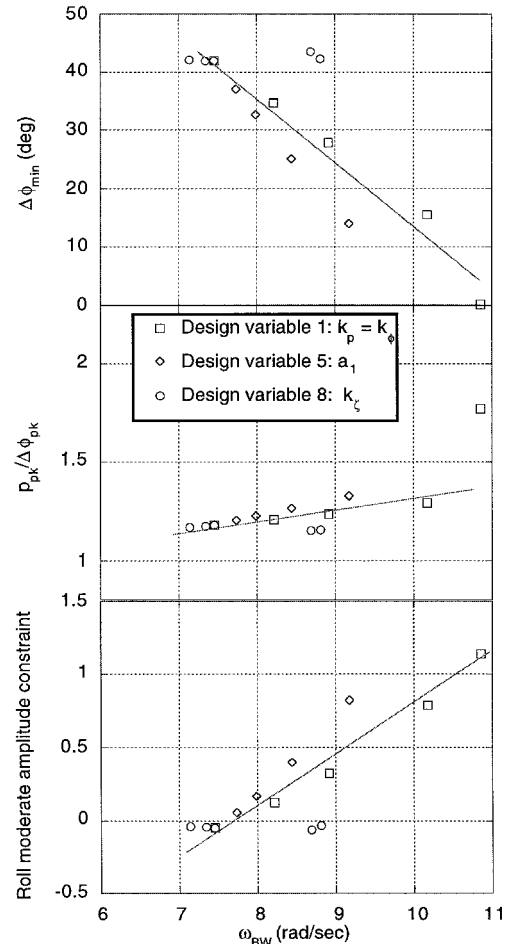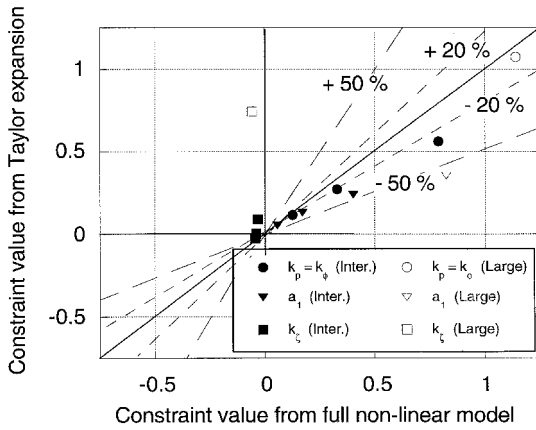

Fig. 10   Roll moderate amplitude constraint for different designs.

Fig. 11 Comparison of exact and approximate values of roll moderate amplitude constraint; $l_1(X) = L_{\theta_{1c}}/L_p$.
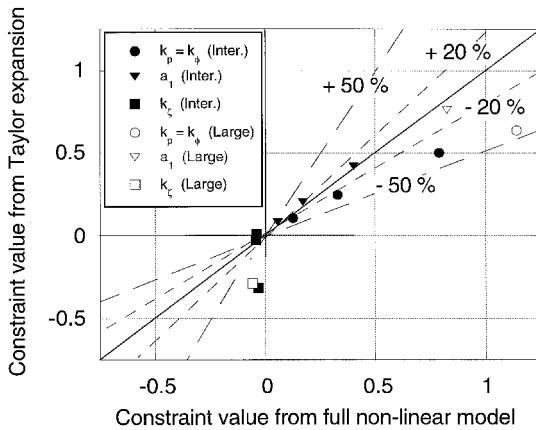

Fig. 12 Comparison of exact and approximate values of roll moderate amplitude constraint; intermediate variable $l_2(X) = \omega_{BW}$.

linear Taylor-series expansions in the intermediate design variables l_1 and l_2 can be accurate for larger moves away from the current design. This is confirmed by Figs. 11 and 12 that compare the exact values of the quickness constraint with those obtained from linear Taylor-series expansions in l_1 and l_2, respectively. Figure 11 shows that when $L_{\theta_{1c}}/L_p$ is used as intermediate variable and $X_1 = k_p = k_\phi$ or $X_5 = a_1$ are perturbed by even large amounts, the expansion predicts the correct sign of the constraint and the quantitative agreement is reasonable. This is not true for changes in $X_8 = k_\zeta$, which mostly affects the rotor lag frequency. Substantial violations are predicted when instead the constraint is satisfied. Because in the present study this Taylor series was not used in an actual optimization, the precise effect of this inaccuracy cannot be determined. However, since rotor lag frequency is not likely to affect significantly the roll (or pitch) moderate amplitude constraint, X_8 might simply be dropped from the vector of design variables when evaluating this constraint. No such problems occur when the intermediate variable is the bandwidth ω_{BW}, as shown in Fig. 12: the Taylor-series expansion predicts the correct sign of the constraint and the quantitative agreement is reasonably good.

Conclusions

This paper describes a new technique for the efficient calculation of gradients of constraints associated with the moderate amplitude attitude change criterion of the ADS-33 handling qualities specifications. The gradients are calculated using low-order linear approximations to the full nonlinear mathematical model of the helicopter. No modifications to existing flight dynamics simulation codes are required to make

use of this technique. The sensitivities of the full-order linearized model with respect to the design variables are assumed to be available. This includes the sensitivities of state matrix A, control matrix B, and bandwidth ω_{BW}. A limitation of this study is that actuator saturation is not included in the simulation model.

The main results of this study are as follows:

1) Low-order linear models can be obtained either in transfer function form through a curve fit in the frequency domain, or in state-space form by condensation. Both types can provide good approximations to the gradients of the moderate amplitude constraint. These models are not sufficiently accurate for the calculation of the perturbed constraint, which should be obtained by adding the precise baseline value of the constraint to the increment calculated using the approximate gradients.

2) The technique presented in this paper reduces, by a factor of about 50, the additional cost of calculating the gradient information once the baseline value is available, compared with traditional one-sided finite differences.

3) When used in an actual optimization, these approximate gradients make it achieve most of the reduction in the objective function that could be obtained using the full nonlinear model and do not cause any additional infeasible intermediate designs.

4) The accuracy of linear Taylor-series expansions of the constraint depends on the size of the changes of each design variable; a change of 40% or less will usually give reasonable accuracy. There is no improvement in accuracy if the expansions are in terms of intermediate variables, such as reciprocals and cubes of the design variables. The accuracy improves if the bandwidth or the derivative ratio $L_{\theta_{1c}}/L_p$ are used as intermediate variables, in which case perturbations of 80% (or possibly larger) can become acceptable.

References

[1]Sahasrabudhe, V., Celi, R., and Tits, A. L., "Integrated Rotor-Flight Control System Optimization with Aeroelastic and Handling Qualities Constraints," *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 2, 1997, pp. 217–224.

[2]Anon., "Aeronautical Design Standard ADS-33D, Handling Qualities Requirements for Military Rotorcraft," U.S. Army Aviation and Troop Command, St. Louis, MO, July 1994.

[3]Tomović, R., *Sensitivity Analysis of Dynamic Systems*, McGraw–Hill, New York, 1963.

[4]Wilkie, D. F., and Perkins, W. R., "Generation of Sensitivity Functions for Linear Systems Using Low-Order Models," *IEEE Transactions on Automatic Controls*, Vol. AC-14, No. 2, 1969, pp. 123–130.

[5]Adelman, H. M., and Haftka, R. T., "Sensitivity Analysis of Discrete Structural Systems," *AIAA Journal*, Vol. 24, No. 5, 1986, pp. 823–831.

[6]Haftka, R. T., and Gurdal, Z., *Elements of Structural Optimization*, 3rd ed., Kluwer Academic, Norwell, MA, 1992.

[7]Ray, D., Pister, K. S., and Polak, E., "Sensitivity Analysis for Hysteretic Dynamic Systems: Theory and Applications," *Computer Methods in Applied Mechanics and Engineering*, Vol. 14, 1978, pp. 179–208.

[8]Kramer, M. A., Calo, J. M., and Rabitz, H., "An Improved Computational Method for Sensitivity Analysis: Green's Function Method with AIM," *Applied Mathematics and Modeling*, Vol. 5, No. 6, 1981, pp. 432–441.

[9]Feng, T.-T., Arora, J. S., and Haug, E. J., "Optimal Structural Design Under Dynamic Loads," *International Journal for Numerical Methods in Engineering*, Vol. 11, No. 1, 1977, pp. 39–52.

[10]Hsieh, C. C., and Arora, J. S., "Design Sensitivity Analysis and Optimization of Dynamic Response," *Computer Methods in Applied Mechanics and Engineering*, Vol. 43, No. 2, 1984, pp. 195–219.

[11]Grandhi, V. R., Haftka, R. T., and Watson, L. T., "Design-Oriented Identification of Critical Times in Transient Response," *AIAA Journal*, Vol. 24, No. 5, 1986, pp. 649–656.

[12]Barthelemy, J.-F. M., and Haftka, R. T., "Recent Advances in Approximation Concepts for Optimal Structural Design," *Optimization of Large Structural Systems*, Vol. 1, Kluwer Academic, Norwell, MA, 1993, pp. 235–256.

[13]Vanderplaats, G. N., *Numerical Optimization Techniques for Engineering Design: With Applications*, McGraw–Hill, New York, 1984.

[14]Zhou, J. L., and Tits, A. L., "User's Guide for FFSQP Version 3.5," Systems Research Center, Univ. of Maryland, TR-92-107r4, College Park, MD, 1995.

[15]Vanderplaats, G. N., *DOT—Design Optimization Tools, User's Manual*, VMA Engineering, Inc., Goleta, CA, May 1995.

[16]Landis, K. A., and Glusman, S. I., "Development of ADOCS Controllers and Control Laws," NASA CR 177339, 1984.

[17]Ormiston, R. A., and Hodges, D. H., "Linear Flap-Lag Dynamics of Hingeless Helicopter Rotor Blades in Hover," *Journal of the American Helicopter Society*, Vol. 17, No. 2, 1972, pp. 2–14.

[18]Spence, A. M., and Celi, R., "Efficient Sensitivity Analysis for Rotary-Wing Aeromechanical Problems," *AIAA Journal*, Vol. 32, No. 12, 1994, pp. 2237–2344.

[19]Heffley, R. K., "A Compilation of Analysis of Helicopter Handling Qualities Data," NASA CR 3145, Aug. 1979.