

An Analytic Modeling And System Identification Study of Rotor/Fuselage Dynamics At Hover

Steven W. Hong
United Technologies Research Center
East Hartford, Connecticut

H. C. Curtiss, Jr., Professor
Princeton University
Princeton, New Jersey

Abstract

A combination of analytic modeling and system identification methods have been used to develop an improved dynamic model describing the response of articulated rotor helicopters to control inputs. A high-order linearized model of coupled rotor/body dynamics including flap and lag degrees of freedom and inflow dynamics with literal coefficients is compared to flight test data from single rotor helicopters in the near hover trim condition. The identification problem was formulated using the maximum likelihood function in the time domain. The dynamic model with literal coefficients was used to generate the model states, and the model was parametrized in terms of physical constants of the aircraft rather than the stability derivatives, resulting in a significant reduction in the number of quantities to be identified. The likelihood function was optimized using the genetic algorithm approach. This method proved highly effective in producing an estimated model from flight test data which included coupled fuselage/rotor dynamics. Using this approach it has been shown that blade flexibility is a significant contributing factor to the discrepancies between theory and experiment shown in previous studies. Addition of flexible modes, properly incorporating the constraint due to the lag dampers, results in excellent agreement between flight test and theory, especially in the high frequency range.

Presented at Piloting Vertical Flight Aircraft: A Conference On Flying Qualities and Human Factors, San Francisco, California, 1993.

Introduction

The investigation of rotorcraft dynamics, and specifically the coupled fuselage/rotor dynamics, is motivated by increasing sophistication in rotorcraft stability analyses and by the emergence of high-performance flight control system design requirements. The past few years have seen a concentrated effort directed toward providing an analytic simulation model of coupled fuselage/rotor dynamics and model validation against flight test data.

Helicopter dynamics include the rigid-body responses demonstrated by fixed-wing aircraft, plus higher-frequency modes generated by the interactions of the rotor system with the fuselage. For earlier flight control system designs with lower bandwidth requirements, it was satisfactory to use low-order analytic models which did not accurately model the high-frequency rotor dynamics; with the recent introduction of high-performance, high-bandwidth control system specifications, it has become increasingly necessary to correctly model the coupled fuselage/rotor dynamic modes. It has long been known that flap dynamics introduce significant time delays into the rotor system, and more recently, Curtiss has shown that inclusion of the lag dynamics is important in the design of high performance control systems (Curtiss, 1986). Recent studies have explored the possibility of using rotor state feedback designs to damp blade motion (Ham, 1983). An accurate understanding of the coupled fuselage/rotor dynamics is therefore important in rotorcraft control system design and stability analyses.

Recent flight test experiments have shown that existing simulation models do not accurately

predict these high-frequency modes (Ballin et. al, 1991, Kaplita et. al, 1989, and Kim et. al, 1990). These studies show significant differences between theory and experiment associated with the coupled rotor/body dynamics, especially in the frequency region dominated by the rotor lag motion. This research is therefore directed toward providing an improved understanding of the aeroelastic and aeromechanical phenomena which determine the coupled rotor/body dynamics at hover.

In order to gain physical insight into helicopter dynamics, development of linear models incorporating coupled rotor/fuselage dynamics has long been a research objective. Past approaches to linear model development have included direct numerical perturbation of nonlinear simulations (Diftler, 1988), identification of state-space stability and control matrix elements (Tischler, 1987), and analytic derivation of linear equations of motion (Zhao and Curtiss, 1988). This study uniquely combines system identification methods with analytic modeling techniques in order to investigate helicopter hover dynamics and to arrive at an improved linear model. The emphasis is on the high-frequency dynamics of the coupled rotor/body motion.

The identification study is carried out on flight test data from a Sikorsky H-53E helicopter at hover, using previously published data (Kaplita et. al, 1987, and Mayo et. al, 1990).

Research Objectives

This paper describes an investigation into the response of articulated rotor helicopters to control inputs in hover. The goal is an improved understanding of the coupled rotor/fuselage dynamics in hover directed toward a validated analytic simulation model including high-frequency rotor/fuselage dynamics for use in stability analyses and high-performance control system design studies.

Identification of linear, time-invariant state-space models representing high-order helicopter dynamics including main rotor degrees of freedom has long been an objective of engineers involved in rotorcraft simulation and control system design. The state and control matrix elements in an identified state-space model can provide physical insight

into system dynamics and can be used in combination with mathematical modeling techniques to analyze differences between theory and experiment.

State-space identification techniques have been applied to conventional fixed-wing aircraft with useful results. Since identification of state-space models using directly parametrized state and control matrix elements requires the estimation of a large number of parameters, a reduced order model is often used, assuming six degree-of-freedom rigid body dynamics and decoupling between the longitudinal and lateral axes.

Identification of reduced order state-space models for rotorcraft have generally produced unsatisfactory results. The presence of the rotor produces significant rotor/body coupling, requiring additional states to describe the high-frequency dynamics, and also introduces significant interaxis coupling. The complete rotorcraft identification problem is therefore required to use a high-order, multi-input, multi-output model with as many as 18 or more states.

In order to avoid the inevitable problem of overparametrization which results when attempting to identify a directly parametrized high-order helicopter model, this study uses an analytic model to generate state time histories. The model used in this study has been developed at Princeton using the Lagrangian formulation. It includes the coupled fuselage/rotor dynamics, main rotor inflow, tail rotor thrust, and provides for tail rotor inflow dynamics. It was analytically linearized about hover. This model provides a state-space description of the helicopter at hover which is completely analytic and dependent only on an input set of physical parameters. A subset of these inputs are considered uncertain, and are to be estimated from flight test data. The flight-test derived parameter estimates can be used in combination with the mathematical formulation to trace various physical aspects of coupled rotor/body dynamics and thereby obtain physical insight. The complete high-order model including rotor dynamics can be reasonably parametrized by 15 or fewer physically meaningful input coefficients, resulting in a substantial reduction in the number of parameters to be estimated.

The framework of the identification approach is the time-domain maximum likelihood methodology. The likelihood function is formulated assuming the presence of Gaussian measurement and process noise. The process noise may be nonwhite. The noise covariances as well as process noise dynamics may be parametrized. With Gaussian noise assumptions, the likelihood function becomes the weighted least-square of the residual errors. The Kalman filter is the natural way to produce these residuals for state-space dynamic systems.

The maximum likelihood estimate is obtained by finding the global maximum of the likelihood function. The parameters are nonlinearly related to the cost function and the resulting parameter space is highly multimodal. Traditional function optimization techniques based on gradient methods generally become trapped in local optima.

The genetic algorithm is an alternative function optimization approach which does not rely on the use of local gradient information. The genetic algorithm is an adaptive scheme, based on the analogy with natural evolution, which efficiently searches a large parameter space for the 'fittest' solution to a given objective. This method has been demonstrated to be highly effective in obtaining the global maximum in a multimodal parameter space.

The formulation of the system identification problem in the maximum likelihood framework leads to estimates of physical coefficients which have attractive statistical optimality properties and represent the best possible combination of physical coefficients necessary to match the given test data set.

This identification methodology allows an assessment of model assumptions inherent in the mathematical model used to generate the state time histories. In this study, emphasis is placed on the frequency region associated with coupled rotor/fuselage dynamics. In the frequency domain, the dominant feature in the rotor magnitude response is a notch characteristic produced by the presence of the in-plane blade degree of freedom. Using rotor blade constants derived through the identification procedure, rotor blade modeling assumptions may be examined, resulting in analytic model improve-

ments. This study examines in detail the blade structural modeling assumption and investigates the effect of accounting for blade flexibility effects generated by the presence of a large mechanical damper at the blade hinge.

Analytic Model Description

Research at Princeton has resulted in the development of a linearized rotor/body helicopter dynamic model. The dynamic equations are formulated using a Lagrangian approach in order to capture all the important inertial coupling terms. The model includes rigid-body translation and rotation (pitch, roll, and yaw rates, longitudinal and lateral velocities), rigid blade lag and flap multimodal coordinates, and main rotor cyclic dynamic inflow. The controls are main rotor cyclic and pedals. The version of the model used in this study was analytically linearized about the hover trim condition and does not include the collective degree of freedom.

Rotorcraft dynamics includes coupling between the motion of the fuselage which is in rotational and translational motion relative to inertial space, and the motion of individual rotor blades. The final set of equations of motion are referenced to the body-fixed axis system which has its origin at the fuselage center of gravity. In the Newtonian approach to modeling coupled rotor/fuselage equations of motion, blade acceleration terms are first written referenced to the hub axis which is rotating at constant velocity; coordinate transformations are then used to obtain acceleration terms in the body-fixed frame. The complexity of the resulting acceleration terms, combined with the number of degrees of freedom necessary to model rotor dynamics properly, has led to the use of Lagrange's equations for the derivation of the coupled rotor/body model.

The development of Lagrange's equations proceeds from the evaluation of the Lagrangian, which requires only position and velocity terms in order to relate the system generalized forces to changes in the system kinetic and potential energies. The generalized coordinates in Lagrange's approach represent the degrees of freedom in the system and are chosen to correspond to the system

states. The kinetic energy term includes the motion of the fuselage and rotor blades, and the potential energy includes the gravitational potential energy of the fuselage and stored energy in the mechanical springs in the rotor system. Mechanical dampers are accounted for by use of the dissipation function. The generalized forces include aerodynamic forces due to fuselage and blade aerodynamics. Evaluation of the time and partial derivatives in the Lagrangian can be time consuming for a high-order model and can be assigned to a symbolic manipulation program such as MACSYMA.

Identification Methodology

This paper describes an approach for identification of a coupled fuselage/rotor model for rotorcraft hover dynamics from flight test measurements. The identified model includes flap and lag degrees of freedom, main rotor inflow, and process and measurement noise disturbances. The process noise may be colored. The approach uses an analytically derived, linear time-invariant state-space model with literal coefficients which is parametrized in terms of aeromechanical input coefficients. The model order and structure may therefore be assumed to be determined by this approach, and the system parameters are to be estimated from observations. The parameter estimation problem is formulated using the statistical framework of maximum likelihood (ML) estimation theory, thereby benefitting from known optimality properties of ML estimators. This discussion first presents the parametrized dynamic model to be used in the identification methodology, and then describes the application of the maximum likelihood estimation approach to dynamic systems.

Model Parametrization

The helicopter is modeled as a continuous-time dynamic system whose measurements are discretely sampled as sensor outputs. Thus the identification algorithm is required to estimate continuous-time model parameters from discrete sensor measurements. This continuous/discrete formulation is well known and is discussed by

Ljung (1987). The linear time-invariant state equations are derived using the Lagrangian approach, and are given by

$$\dot{x}(t) = A_c(\theta)x(t) + B_c(\theta)u(t) + F_c(\theta)w(t) \quad (1)$$

The model form accounts for the presence of process noise, where $w(t)$ is assumed to be zero-mean white noise with unity spectral density. The continuous-time matrices, $A_c(\theta)$, $B_c(\theta)$, and $F_c(\theta)$, are parametrized by a vector of parameters, θ , which are to be estimated from observations.

The observations are sampled at discrete time intervals, where

$$y(kT) = C(\theta)x(kT) + G(\theta)v_r(kT) \\ t = kT, \quad k = 0, 1, 2, \dots \quad (2)$$

and $v_r(kT)$ are the disturbance effects at the sampled time intervals.

For digital implementation of the identification algorithm, the continuous-time state equation given in Equation (1) is discretized using zero-order hold. The input is assumed to be held constant over the sampling time interval, and the continuous-time state equation can then be integrated analytically over the interval in order to obtain the discrete-time state equation. The zero-order hold discretization introduces a phase lag equivalent to one-half sample interval, which is taken into account by advancing the control input by the corresponding one-half time interval.

Eliminating time subscripts for simplicity, the discrete-time state-space equations are given by

$$x(t+1) = A(\theta)x(t) + B(\theta)u(t) + F(\theta)w(t) \\ y(t) = C(\theta)x(t) + G(\theta)v(t) \quad (3)$$

This equation is now understood to be a discrete-time equation. Here, $w(t)$ and $v(t)$ are sequences of independent random variables with zero mean and unit covariance.

Maximum Likelihood Formulation

Let Y^N be a vector of observations which are supposed to be realizations of stochastic variables,

and let $y(t)$ be a multi-dimensional observation taken at time t :

$$Y^N = [y(1), y(2), \dots, y(N)]$$

The observations, Y^N , depend on a vector of parameters, θ , which are also considered to be random variables. The conditional probability density function for θ , given the observations, Y^N , is then given by

$$p(\theta|Y^N) = \frac{p(Y^N|\theta) \cdot p(\theta)}{p(Y^N)} \quad (4)$$

where $p(\theta)$ is the prior distribution of the random parameter vector. A reasonable estimate for θ can then be obtained by finding the value of θ which maximizes the conditional density function given by Equation (4). With no prior knowledge of the distribution of θ , $p(\theta)$ may be assumed to be uniform. The best estimate for θ is then obtained by maximizing the likelihood of obtaining the observations. This leads to the ML, or maximum likelihood, estimator, given by

$$\hat{\theta}_{ML} = \arg \max_{\theta} p(Y^N|\theta) \quad (5)$$

For parametrized dynamical systems, with Gaussian noise assumptions, the maximum likelihood estimator has the form

$$\begin{aligned} \hat{\theta}_{ML} &= \arg \max_{\theta} p(Y^N|\theta) \\ &= \arg \max_{\theta} -\frac{1}{2} \sum_{t=1}^N \varepsilon^T(t, \theta) \Lambda^{-1}(\theta) \varepsilon(t, \theta) - \\ &\quad \frac{N}{2} \log |\Lambda(\theta)| - \frac{Nm}{2} \log 2\pi \end{aligned} \quad (6)$$

where

m = number of measurements

$$\varepsilon(t, \theta) = y(t) - \hat{y}(t, \theta)$$

$$\Lambda(\theta) = E\varepsilon(\theta)\varepsilon^T(\theta)$$

and $\hat{y}(t, \theta)$ is generated using Equation (3) with the discrete-time Kalman filter formulation.

The Genetic Algorithm

The evaluation of the likelihood function as presented in Equation (6) requires a search for the global maximum of the likelihood function over a multimodal parameter space whose contours are not known. Specifically, the identification methodology has led to a function optimization problem where the performance measure is a highly nonlinear function of many parameters. The principal challenge facing the identification problem is the very large set of possible solutions and the presence of many local optima. Hill-climbing methods for function optimization based on finding local gradients become trapped in local optima and are inadequate for this problem. Genetic algorithms overcome these difficulties by efficiently searching the parameter space while preserving and incorporating the best characteristics as the search progresses.

The problem of function optimization can be addressed using the paradigm of adaptive systems, where some objective performance measure (the cost function) is to be maximized (i.e., adaptation occurs) in a partially known and perhaps changing environment. The idea of artificial adaptive plans, based on an analogy with genetic evolution, was formally described by John Holland in the seventies and have recently become an important tool in function optimization and machine learning (Holland, 1975, and Goldberg, 1989). Holland's artificial adaptive plans have come to be known in recent literature as genetic algorithms.

Genetic algorithms are based on ideas underlying the process of evolution; i.e., natural selection and survival of the fittest. Using biological evolution as an analogy, genetic algorithms maintain a population of candidate solutions, or 'individuals,' whose characteristics evolve according to specific genetic operations in order to solve a given task in an optimal way.

As a general overview, genetic algorithms have the following attributes which distinguish them from traditional hill-climbing optimization methods (Goldberg, 1989):

1. GA's work with a representation of the parameter values rather than with the parameters themselves.
2. GA's search from a population of points, not from a single point.
3. GA's use objective function information, not gradient information.
4. GA's use probabilistic transition rules, not deterministic ones.

The genetic algorithm maintains a population of 'individuals'; i.e., possible solutions to the function optimization problem. In the context of the identification problem, each individual corresponds to a vector of parameters. The population of individuals 'evolves' according to the rules of reproduction and mutation analogous to those found in natural evolutionary processes, with the result that the population preserves those characteristics favoring the best solution to the cost function.

The following steps were described by Holland (Holland, 1975) and contain the essentials properties of the the basic genetic algorithm.

1. Select one individual from the initial population probabilistically, after assigning each individual a probability proportional to its observed performance.
2. Copy the selected individual, then apply genetic operators to the copy to produce a new individual.
3. Select a second individual from the population at random (all elements equally likely) and replace it by the new individual produced in step 2.
4. Observe and record the performance of the new structure.
5. Return to step 1.

This deceptively simple set of instructions contains the ability to test large numbers of new combinations of individual characteristics and the ability to progressively exploit the best observed characteristics. It does so through the use of genetic operators.

Genetic Operators

Parent selection based on fitness, and the subsequent application of genetic operators to produce new individuals are the steps by which the algorithm modifies the initial population and continually tests new combinations while maintaining those parameter sets which give high fitness. Each of these operations are performed probabilistically.

The initial population of individuals is selected randomly with a uniform distribution over the defined parameter space. After one generation, parent individuals are selected randomly, with a probability which is proportional to the fitness assigned to that individual. The selection procedure resembles spinning a roulette wheel whose circumference is divided into as many segments as there are individuals. The arc length of each segment is made proportional to the fitness value of the corresponding individual. Thus, the chance of choosing a given individual is uniformly random and yet proportional to its fitness.

The genetic operations of crossover and mutation are then applied to the selected parent individuals in order to introduce new characteristics into the population, enabling an efficient search for the optimal combination of parameters.

The crossover operation involves a recombination of two selected individuals at a randomly selected point. Thus the crossover operation produces two new individuals, each of whom inherit characteristics from both parents.

The mutation operation involves a random alternation of an individual's characteristic with a very low probability. This serves to introduce new information into the pool of structures and serves to guard against the possibility of becoming trapped in local optima.

Genetic Coding

Each individual is a candidate parameter set and is represented as a concatenation of individual parameters:

$$\theta = [\theta_1, \theta_2, \dots, \theta_N]$$

In a digital implementation, each parameter θ_i is encoded using a binary alphabet, and the individual is thus represented by a binary-valued string. The following specific coding scheme was suggested by Storer (Storer, 1990).

Let each parameter θ_i be bounded by $\theta_{i_{max}}$ and $\theta_{i_{min}}$. If each parameter is coded in binary with a word length of l , then the interval $[\theta_{i_{max}}, \theta_{i_{min}}]$ is discretized by 2^l values. A representation of the parameter θ_i can be obtained from the l -bit binary coding of

$$\text{mod} \left[\frac{(\theta_i - \theta_{i_{min}})(2^l - 1)}{\theta_{i_{max}} - \theta_{i_{min}}} \right]$$

To illustrate, let an individual represent a candidate parametrization where

$$\theta = [\theta_1, \theta_2] = [3 \quad 4.5]$$

and bounds are given as

$$1 < \theta_1 < 4, \quad 2 < \theta_2 < 7, \quad l = 6$$

The binary-valued string representing this candidate vector is then

$$\theta_{binary} = [101010 \ 011111]$$

The genetic algorithm is illustrated in Figure

1.

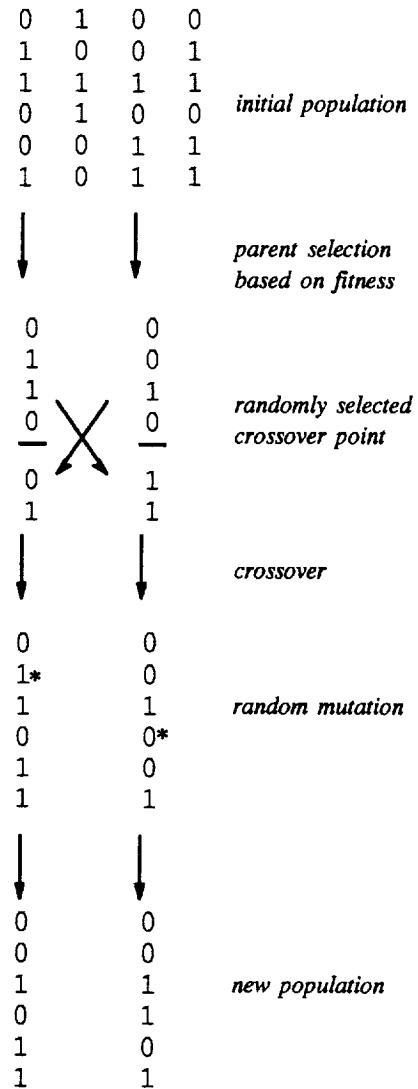


Figure 1 The Genetic Algorithm

Implicit Parallelism

Genetic algorithms efficiently conduct a search over a defined parameter space, converging

to a near-optimal solution. The basic unit of processed information in this genetic search is the schema, defined by Holland (1975). In the context of a digital implementation of genetic algorithms, a schema is a template specifying similarities at certain string positions.

Thus, an individual is a string of binary digits, and the alphabet is composed of $\{0, 1, \#\}$, where # denotes 'don't care' (i.e., the value at this position has no effect on the performance measurement). As an example, an individual may be represented as

[0 0 1 1 1 0 1 1 0 0 0 1 0]

A schema is a similarity template within this individual; so that this individual contains the schemata given by

[0 0 # # 1 0 1 1 0 0 0 1 0]

Given l positions, a single individual is an instance of 2^l distinct combinations, and an instance of 3^l distinct schemata. Further, a population of size N contains between 3^l and $N3^l$ distinct schemata. Holland has shown that each schemata are evaluated and processed independently of the others, providing a tremendous computational leverage on the number of function evaluations. Therefore, the use of genetic operators in the reproductive plan provides i) intrinsic parallelism in the testing and use of many schemata, and ii) compact storage and use of large amounts of information resulting from prior observations of schemata.

The concept of implicit parallelism is fundamental to the efficiency of genetic algorithms. Each schemata is processed and evaluated independently of other schema in the population; this provides a tremendous computational leverage. A very weak lower bound states that for a population of (n) individuals, more than $o(n^3)$ useful 'pieces' of information is processed in each iteration (Goldberg, 1989).

An Example

As an illustration of the genetic algorithm, consider the following example.

$$f(x, y) = 3(1 - y)^2 e^{-y^2 - (x+1)^2} - 10\left(\frac{y}{5} - y^3 - x^5\right) e^{-y^2 - x^2} - \frac{1}{3} e^{-(y+1)^2 - x^2}$$

The function surface is shown in Figure 2, along with the contour lines. This multimodal function has a global maximum at $(1.5814, -0.0093)$.

A genetic algorithm was run on this function with a population size of 20. The initial guesses were chosen randomly, and were bounded as $-3 < x < 3, -3 < y < 3$. A binary code with wordlength of 8 was used, which means that both x and y were discretized by 256 points. An exhaustive grid search under these conditions would involve evaluating 65536 possible points to find the global maximum.


Snapshots of the population distribution up to 7 generations are shown in Figure 2. The snapshots show the population converging upon the global maximum; by the 7th generation, most of the individuals have converged on the maximum. The genetic algorithm in this case converges on $(1.5412, -0.0353)$ as the global optimum.

This convergence has occurred after 7 generations. With a population size of 20 individuals, this is 140 function evaluations as compared to the 65536 necessary for the grid search.


This relatively simple example serves to illustrate the ability of the genetic algorithm to find the optimum of a given function, *using no gradient information*.

Analytic Model Validation


The mathematical model is correlated with flight test data using nominal values for input coefficients. The correlation plots in Figure 3 show transfer function comparisons for pitch and roll axes. The data represent separate flights. In each




Sample Function




Generation 1



Generation 3



Generation 5



Generation 7

Figure 2 Genetic Algorithm Example

case, the comparison is between the flight test rate gyro output and the model state. The comparison is made between 0.5 Hz (3.14 rad/sec) and 6 Hz (37.7 rad/sec) since the input signal was designed to cover this frequency range. The fuselage structural bending modes are lightly damped and dominate the frequency above ~20 rad/sec. Therefore the identification procedure uses a bandpass filter with the upper cutoff frequency at 15.7 rad/sec. The frequency range of interest is therefore between 0.5 Hz to 2.5 Hz (3.14 rad/sec to 15.7 rad/sec).

The choice of physical coefficients used to parametrize dynamic model must allow adjustments to account for differences between test and theoretical responses using nominal physical input values. The gain differences at low frequencies, implying a mismatch in rigid body response, requires parametrization of the rigid body acceleration. The coupled fuselage/lagwise modes are a lightly damped pole-zero pair and create a notch-filter effect in the frequency response between 10 - 15 rad/sec. This frequency is near the -180 degree crossover, and a mismatch in this region adversely impacts the gain and phase margin calculations. Modeling the dynamics of this mode is important for control system design and stability analysis and will be the primary focus of modeling in this study.

Validation Of Identification Procedure Using Simulated Data

The maximum likelihood identification methodology for parametrized dynamic systems is validated first on a simulation with known parameters. These results demonstrate the feasibility of using genetic algorithms to estimate physical coefficients from noisy data, and establish the population size and crossover and mutation rates for this application.

The simulation model is driven by flight test control inputs from the hovering condition. Main rotor pitch and roll cyclic and tail rotor pedals are all active, with primary excitation into roll cyclic. The output states used to form the cost function are pitch, roll, and yaw rates, and pitch and roll attitudes. No velocity information is necessary.

Simulation Model Parametrization

The model structure and parametrization was presented in Equations (1) through (3). The continuous-time state space model is analytically derived using the Lagrangian approach and using a vector of physical input coefficients, θ . For the purposes of this simulation study, the model structure has been augmented to include a first order time constant on process noise. The process noise dynamics are to be parametrized and estimated from output data.

The simulation model was parametrized as follows:

aerodynamic coefficients:

lift curve slope, a
 inflow equivalent cylinder height, $hhnd$
 inflow wake rigidity factor, wrf

hover trim values:

trim flap angle, β_o
 trim main rotor pitch angle, t_o
 trim inflow velocity, v_o

main rotor blade constants:

lag damper constant, C_c
 lag spring constant, K_c
 flap spring constant, K_β

inertias:

fuselage cross-moment, I_{xz}

tail rotor:

tail rotor thrust scale factor, K_{TR}

noise parameters:

noise covariance ratio, NR
 process noise time constant, τ

Kalman filter theory allows optimal state estimates to be obtained in the presence of state and measurement noise, where the Kalman gain is uniquely determined up to the ratio of process to measurement noise. The noise covariance estimate is therefore parametrized by the ratio of process to measurement noise.

Genetic Algorithm Procedure

The genetic algorithm was implemented using a population size of 500 individuals; a crossover

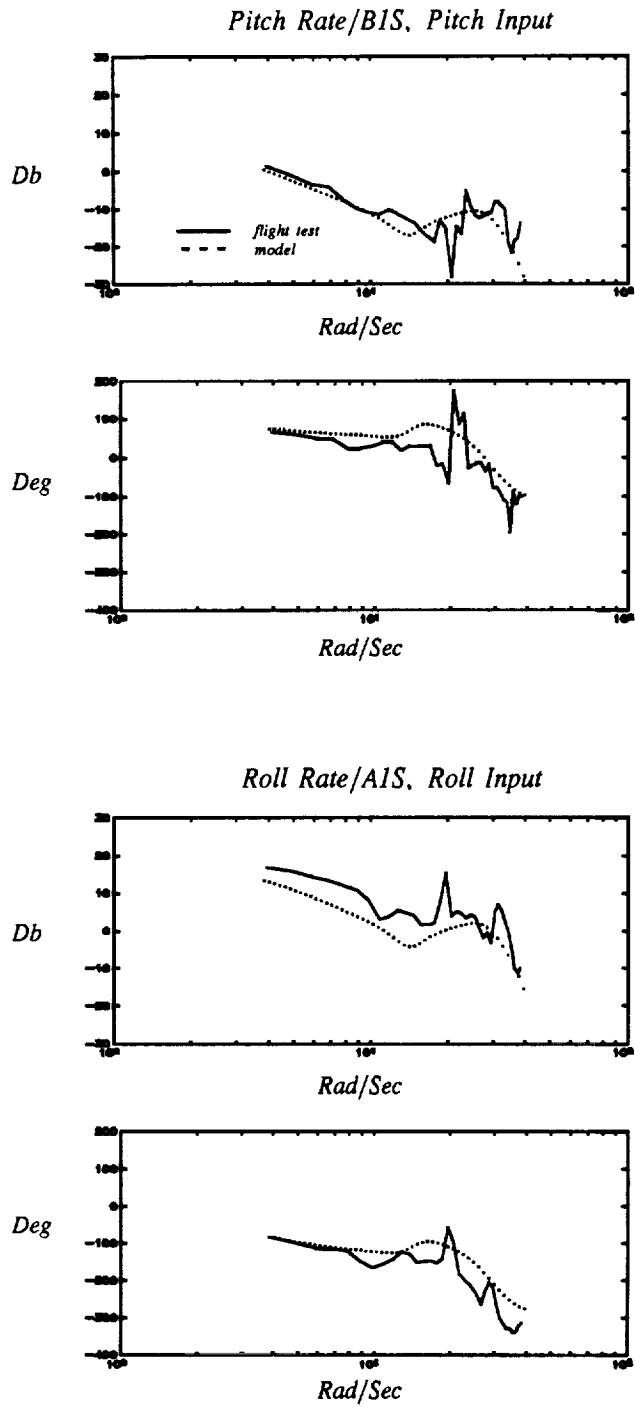


Figure 3 Model Validation

rate of 2/3; and a mutation rate of 1/1000. The parameters were allowed to vary within 50 percent of the known simulation values.




Figure 4 Best Likelihood Values

The sensitivity of the cost function to the parameter values vary widely. Therefore, as parameters begin to show convergence, the range of allowable values is progressively narrowed in order to demonstrate convergence for all parameters.

The identification proceeds by running 10–12 separate genetic algorithms simultaneously, where each algorithm begins with a new random number generator seed to select the initial guesses. Each set of runs therefore produces a scatter band of near optimal guesses for each parameter. The parameters which influence the cost function most are identified most tightly.

Figure 4 shows the progression of the best fitness values out of the population at each generation. The results are shown in Figure 5. The solid line in each figure denotes the true value.

The noise covariance ratio parameter couples only very weakly to the cost function and displays an almost random distribution until the physical coefficient estimates sufficiently converge. Therefore a two-step estimation procedure is required, where the noise ratio is allowed to remain free until physical coefficients have converged. The physical coefficients are then fixed while the noise ratio is estimated.

This methodology clearly demonstrates convergence. Twenty iterations of the genetic algorithm were run. Table 1 tabulates the parameter estimates.

Table 1 Estimated Parameters, Simulation Study

Parameters	θ_o	$\hat{\theta}_o$	std
lift curve slope, a	5.73	5.72	3.98e-4
inflow equivalent cylinder height, $hhnd$	0.46	0.46	2.23e-4
inflow wake rigidity factor, wrf	2.0	2.0	1.34e-4
trim flap angle, β_o	0.02	0.02	4.99e-7
trim main rotor pitch angle, t_o	0.05	0.0497	9.75e-6
trim inflow velocity, v_o	0.02	0.0196	2.61e-6
lag damper constant, \bar{C}_ζ	5.0	4.978	7.7e-3
lag spring constant, K_ζ	75.0	75.0	7.06e-2
flap spring constant, K_β	45.0	44.92	6.3e-3
fuselage cross-moment of inertia, I_{xz}	30,000	30,035	4.98
tail rotor thrust factor, K_{TR}	1.0	0.99	9.35e-4
covariance ratio, process/measurement, NR	1.0	0.97	0.11
process noise time constant, τ	-1.0	-0.99	1.8e-3

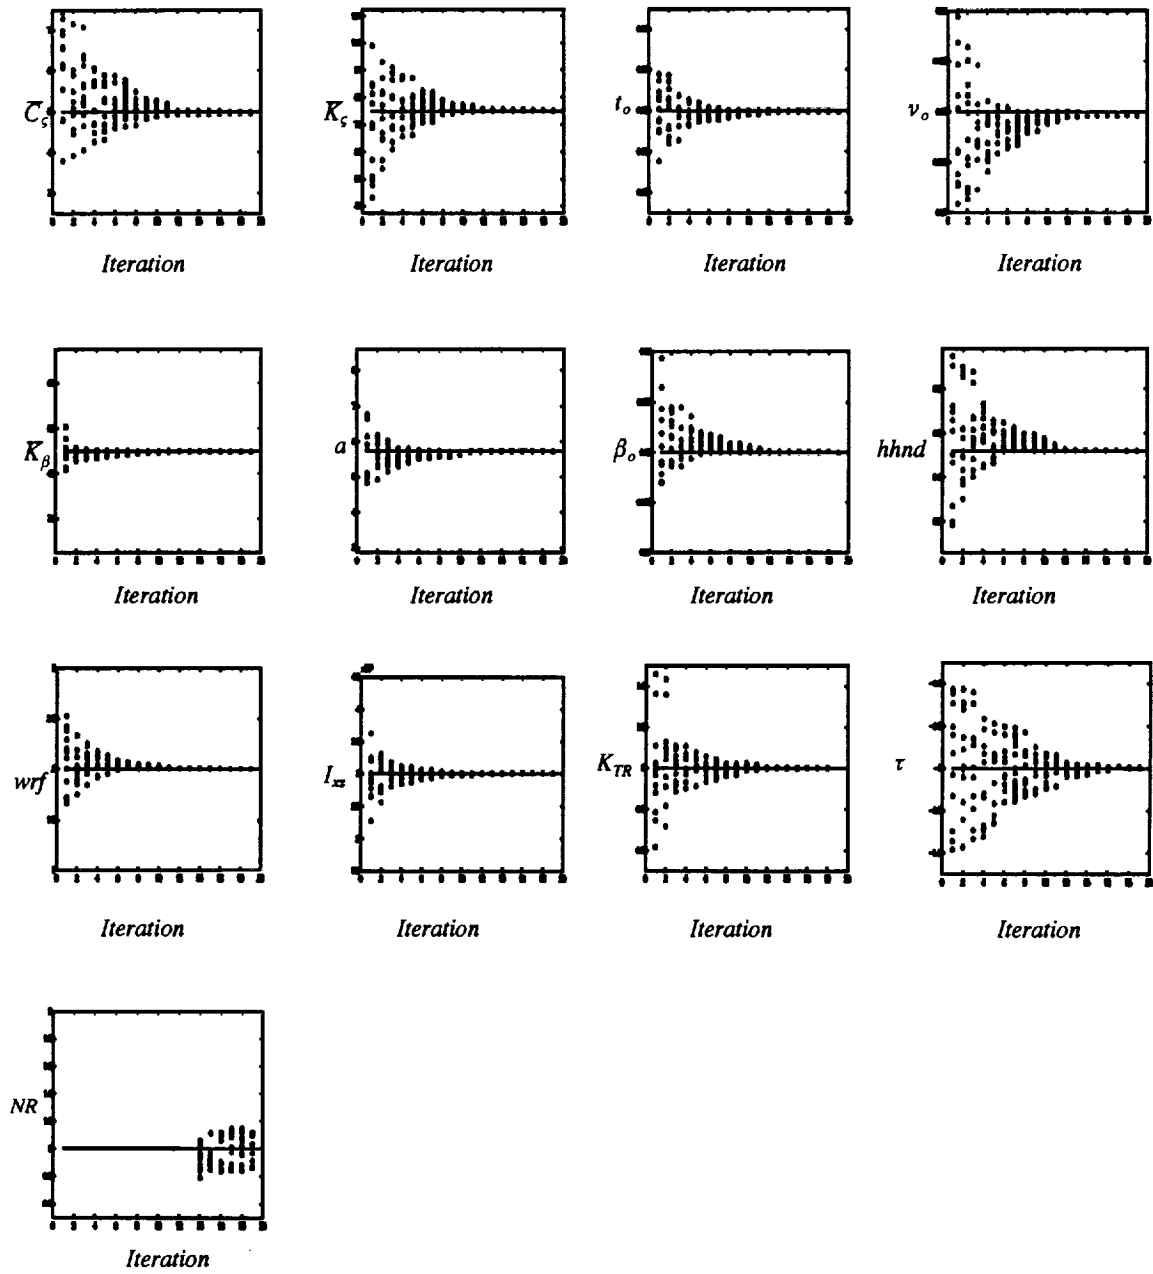


Figure 5 Simulation Identification Results

Flight Test Identification Results

Data consistency checks ensure that errors in data collection do not interfere with the estimation procedure. The requirements for this step were minimal in this study, since this estimation methodology requires only rate and attitude information. Consistency was checked by integrating accelerations and rates, and ensuring that sensor attitudes and rates match the integrated rates and attitudes.

The flight test data was processed by 1) applying a bandpass filter, and 2) decimating the data from 80 Hz to 8 Hz. The filter passband was from 0.5 to 2.5 Hz (3.1416 to 15.708 rad/sec). The lower bound corresponds to the beginning frequency of the frequency sweep input used to drive the system, and the upper bound is imposed to exclude the first fuselage bending mode at 3.4 Hz.

The flight test identification parametrization was modified to reflect information available from comparison between test and theoretical responses generated from the analytic model using nominal parameter values. The parameter list used in flight test identification runs is shown in Table 2. The modifications are explained below.

The parametrization of body inertias accounts for significant differences between theory and test in rigid body response, especially in the roll axis. Further, due to significant differences in

cross-axis predictions, the roll and yaw rigid body responses could not be simultaneously satisfied. Therefore, yaw axis parameters were eliminated, and the identification scheme therefore attempts to fit pitch and roll responses only. This is permissible since for small motions about hover, yaw rate does not couple with main rotor cyclic multiblade coordinates and has no effect on pitch and roll responses in the rotor/body frequency region.

The inflow equivalent cylinder height ($hhnd$) is related to the main rotor dynamic inflow time constant. This parameter had no effect on the cost function in the bandpass frequency region used in this study. Therefore a quasistatic main rotor inflow formulation was used and this parameter was dropped.

The process noise dynamics, parametrized by a first order time constant, was also eliminated. This parameter is uniquely identifiable apart from the noise power ratio only if the time constant falls within the bandpass frequency range, and was found to have no effect on the cost function.

The identification run was carried out using flight test data from hover, with primary excitation into roll cyclic. The analytic model, parametrized as given in Table 2, was driven by main rotor pitch and roll cyclic and tail rotor pedal. The likelihood function was formed using pitch and roll rates only.

Table 2 Estimated Parameters, Flight Test

Parameters	$\hat{\theta}_o$	std	bounds	nominal
scale factor, fuselage roll moment of inertia, I_x	0.44	0.011	0.35-1.0	1.0
scale factor, fuselage pitch moment of inertia, I_y	1.15	0.033	0.7-1.3	1.0
lift curve slope, a	8.4	0.066	5-10	5.73
inflow wake rigidity factor, wrf	8.0	0.23	2-11	2.0
trim flap angle, β_o	0.162	0.0013	0.05-0.25	0.0848
trim main rotor pitch angle, t_o	0.0172	0.00016	0.005-0.15	0.1304
trim inflow velocity, v_o	0.048	0.0007	0.01-0.1	0.0613
lag damper constant, \bar{C}_ζ	5.5	0.10	4-10	9.5
lag spring constant, K_ζ	85.0	0.735	0-100	0
flap spring constant, K_β	16	1.34	0-20	0
noise covariance ratio, NR	-	-	0.001-0.1	-

The initial choice of boundary limits on each parameter defines the parameter space to be searched in the identification algorithm. The bounds applied to each parameter are shown in Table 2; in each case, the bounds are chosen to include the nominal value.

Table 2 shows the identification results for flight test data. It was found that the noise ratio parameter did not converge while the remaining physical coefficients did, indicating that relative to the aeromechanical coefficients, noise powers affect the cost function only very weakly.

The correlation with flight test data using the identified parameters is shown in Figure 6, where the roll axis response is correlated with the data set used in the identification, and the pitch axis response is an independent check. The roll axis correlation shows clear improvement in model correlation using identified coefficients. The low frequency gain prediction has been corrected through the inertia adjustment, and the notch in gain response due to the coupled lag/body response has been corrected.

The differences between identified and nominal parameters can provide physical insight into rotor phenomena when analytic explanations can be found for parameter differences. The identified parameters for lift curve slope, a , and wake rigidity factor, wrf , have produced significant improvement in model response, indicating a possible requirement for refinement of the aerodynamic theory used in the model. The identified parameters for main rotor spring and damping constants indicate necessary refinements in the prediction of frequency and damping of blade motion. A model improvement for blade in-plane dynamics is now presented.

Modeling Blade Elasticity

The identification procedure has resulted in estimated values for rotor blade spring and damping parameters which are different from nominal values. The nominal mechanical damper value may be assumed to be known since it can be independently verified through available data.

A procedure for modeling blade elasticity is presented which accurately accounts for differences between nominal and estimated values for in-plane motion frequency and damping. The method of assumed modes is used to model the case of a flexible beam with damper and spring constraints. This procedure is first demonstrated on a nonrotating beam, for which an exact solution can be obtained. The method of assumed modes will be shown to be a good approximation of the exact solution. This approximate solution can then be used in the flexible beam analysis in the analytic hover helicopter model. The beam formulations for both rotating and nonrotating blades with both spring and damper constraints at the root is given in detail in Appendices A and B.

Approximate solution methods such as the method of assumed modes display convergence toward the analytic solution as more assumed mode shapes are added to the set of basis functions. The first approach to the lagwise bending problem was to use increasing numbers of mode shapes that fulfilled the boundary conditions for a hinged beam. However, with this approach, convergence was not achieved after even after using 5 assumed modes. In order to avoid using an unacceptably large number of basis polynomials in the model, an alternative approach using a combination of modes that satisfy hinged and cantilever boundary conditions was used.

Figure 7 illustrates the assumed modes solution method using both the nonrotating and rotating beam formulations. For a nonrotating beam with spring and damper constraints, an exact expression for the beam eigenvalues is available and is given in detail in Appendix B. The analytic eigenvalue equation is solved numerically. In this case, the root finding problem was converted into a function optimization problem and solved using the genetic algorithm. This solution to the exact formulation is shown against approximate solutions in Figure 7. The approximate solution using the Lagrangian approach, when using only basis functions which fulfill hinged beam boundary conditions, approach the exact solution slowly. With 4 hinged basis polynomials, the solution has not yet converged. However, the assumed modes approach with only one hinged plus one cantilever mode shapes matches

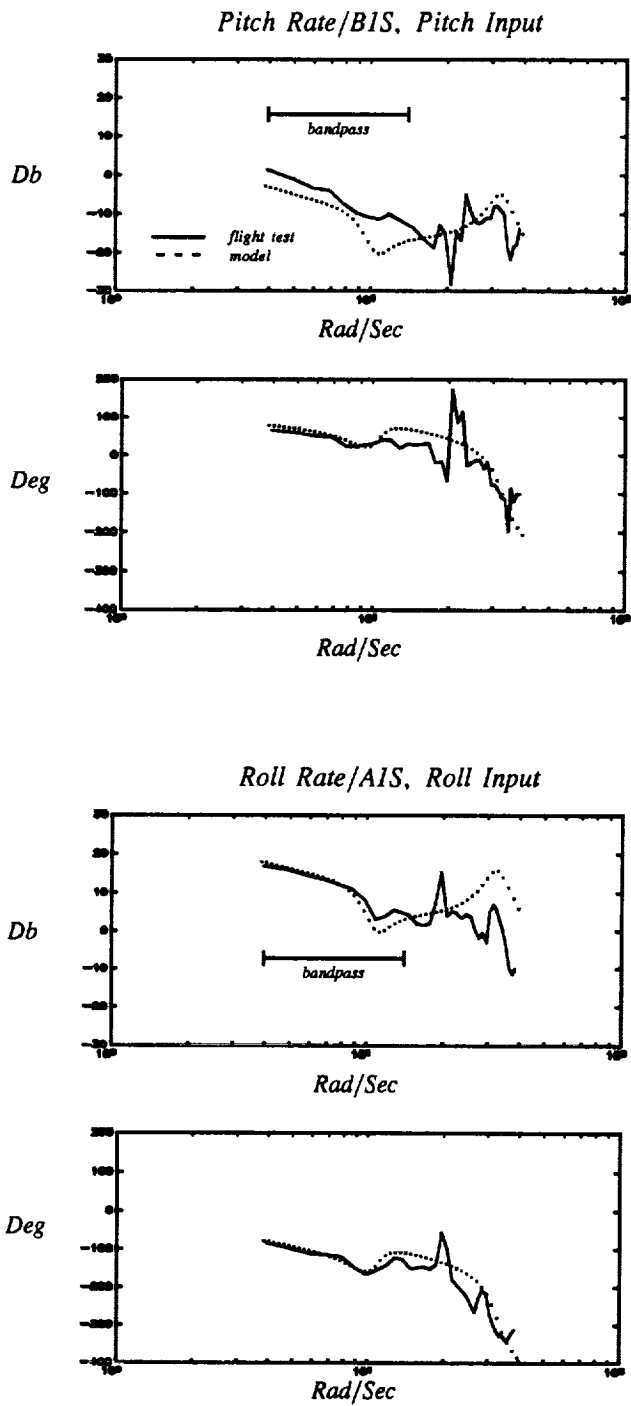


Figure 6 Identified Model Validation





Figure 7 Modal Solutions For Beam Equations

the analytic solution exactly. Convergence is demonstrated by the fact that addition of either hinged or cantilever mode shapes do not further change the eigenvalue solution.

Figure 7 then shows the convergence of the approximate solution for the rotating beam, for which there exists no known exact solution. Here, the sum of 2 hinged plus 2 cantilever modes is near convergence. The addition of either one more hinged or one more cantilever mode does not change the solution appreciably. The combination of 2 hinged plus 2 cantilever modes is chosen for model development as a good compromise between model order and accuracy of solution.




Figure 8 Rotating Frame Lag Roots

Figure 8 shows the location of the rotating frame lag mode eigenvalues. The elastic blade model using two hinged and two cantilever mode shapes is used to show the progression of the root location as damper value is increased from zero to the nominal value. The predicted root location for the elastic model with the nominal damper constant agrees reasonably well with the predicted location for the rigid blade model using a fictitious spring and using identified spring and damper constants. The rigid blade model using nominal damper constant only (no spring) predicts a much higher damping and lower frequency than is indicated by test data.

Conclusions

An analytically derived linear model of coupled rotor/body dynamics at hover has been validated against flight test data.

The analytic model with literal coefficients has been parametrized using 11 physically meaningful coefficients, including noise covariances. This model has been used to formulate a multi-input, multi-output likelihood function in the time domain. The analytic model is used to generate the state time histories. Only body rates are necessary in the cost function.

The likelihood function is globally maximized using the genetic algorithm approach, resulting in statistically optimal maximum likelihood parameter estimates.

The estimated parameters indicate that lag mode damping in flight is approximately one-half of the value expected from rigid blades.

The correct analytic prediction for lagwise motion is obtained using an elastic blade formulation. The flexible blade model was formulated using a normal mode approach and checked using the closed form solution for a nonrotating beam. The convergence results using assumed mode shapes indicate that the correct lagwise bending mode shapes are obtained using a combination of cantilever and hinged assumed modes.

Acknowledgement

One of the authors (S. Hong) would like to thank Mr. Bill Twomey in the Dynamics Section at Sikorsky Aircraft for key discussions on the effect of blade flexibility on blade motion, Professor P. M. Schultheiss at Yale University for insights into the identification framework, and Dr. Richard Williams at the United Technologies Research Center for generously supporting this research effort.

Part of this work was carried out at Princeton University under NASA Ames grant NAG2-561.

Appendix A. Modeling Blade Elasticity

Equation (A.1) gives the in-plane bending equation for a rotating beam. The derivation can be

found in Bramwell (1976), and in Johnson (1980). This partial differential equation relates the moments due to the inertial, centrifugal, and aerodynamic forces to the moment expression from engineering beam theory.

$$\frac{\partial^2}{\partial r^2} \left[EI \frac{\partial^2 Y}{\partial r^2} \right] - \frac{\partial}{\partial r} \left[G(r) \frac{\partial Y}{\partial r} \right] + m \left[\frac{\partial^2 Y}{\partial t^2} - \Omega^2 Y \right] = 0 \quad (A.1)$$

All quantities are understood to refer to lagwise bending motion. Here, $G(r)$ is the centrifugal tension force at a point at a distance r from the hub center, E is the modulus of elasticity, I is the lagwise area moment, and Ω is the rotor rotational velocity.

The boundary conditions for a hinged blade are:

At the hinge:

$$Y(e) = 0$$

$$EI \frac{\partial^2 Y}{\partial r^2} = \text{moment} = 0$$

At the tip:

$$EI \frac{\partial^2 Y}{\partial r^2} = 0$$

$$\frac{\partial^3 Y}{\partial r^3} = \text{shear force} = 0$$

There is no known analytic solution for Equation (A.1) due to the presence of the centrifugal term. A solution based on the method of assumed modes is presented.

Let the lagwise displacement be of the form

$$Y(x, t) = R \sum_n \phi_n(x) q_n(t) \quad (A.2)$$

where R = blade length. This solution method follows the method of separation of variables. $\phi_n(x)$ are a sequence of functions, not necessarily orthogonal, which approximate the expected blade shape and which satisfy the blade boundary conditions.

Substituting into Equation (A.1),

$$\frac{\partial^2}{\partial x^2} EI \sum_n q_n \frac{\partial^2}{\partial x^2} \phi_n - R^2 \frac{\partial}{\partial x} G \sum_n q_n \frac{\partial}{\partial x} \phi_n + \sum_n (\ddot{q}_n - \Omega^2 q_n) \phi_n R^4 m = 0 \quad (A.3)$$

Multiply Equation (A.3) by ϕ_m and integrate from $\frac{e}{R} < x < R$, or $\bar{e} < x < 1$ where \bar{e} is understood to be a nondimensional offset value.

This gives

$$\sum_n q_n \int_{\bar{e}}^1 \phi_m \frac{\partial^2}{\partial x^2} EI \phi_n'' dx - R^2 \sum_n q_n \int_{\bar{e}}^1 \phi_m \frac{\partial}{\partial x} G \phi_n' dx + R^4 \sum_n (\ddot{q}_n - \Omega^2 q_n) \int_{\bar{e}}^1 m \phi_m \phi_n dx = 0 \quad (A.4)$$

Integrating each term by parts, the first term gives

$$\sum_n q_n \int_{\bar{e}}^1 \phi_m \frac{\partial^2}{\partial x^2} EI \phi_n'' dx = \left[\sum_n q_n \phi_m \frac{\partial}{\partial x} EI \phi_n'' \right]_{\bar{e}}^1 - \left[\sum_n q_n \phi_m' EI \phi_n'' \right]_{\bar{e}}^1 + \int_{\bar{e}}^1 \sum_n q_n EI \phi_n'' \phi_m'' dx = RD \sum_n \phi_m' \phi_n' \dot{q}_n + \int_{\bar{e}}^1 \sum_n q_n EI \phi_n'' \phi_m'' dx \quad (A.5)$$

Equation (A.4) was obtained using the boundary conditions for the hinged blade, along with the end constraint imposed by the damper, which is given by

$$EI \frac{\partial^2 Y}{\partial r^2} \Big|_{r=e} = -D \frac{\partial^2 Y}{\partial t \partial r} \Big|_{r=e} = -DR \sum_n \frac{\partial \phi_n}{\partial x} \frac{\partial q_n}{\partial t} \Big|_{x=\frac{e}{R}}$$

where D = damping constant.

Similarly, the second term gives

$$R^2 \sum_n q_n \int_{\bar{e}}^1 \phi_m \frac{\partial}{\partial x} G \phi_n' dx =$$

$$- R^2 \sum_n q_n \int_{\bar{x}}^1 G \phi_n' \phi_m' dx \quad (A.6)$$

Using Equations (A.4) through (A.6),

$$\begin{aligned} & \int_{\bar{x}}^1 \sum_n q_n EI \phi_n'' \phi_m'' dx + RD \sum_n \phi_n' \phi_m' \dot{q}_n \\ & + R^2 \sum_n q_n \int_{\bar{x}}^1 G \phi_n' \phi_m' dx \\ & + R^4 \sum_n (\ddot{q}_n - \Omega^2 q_n) \int_{\bar{x}}^1 m \phi_n \phi_m dx = 0 \end{aligned}$$

To evaluate this, nondimensionalize by $m\Omega^2 R^4$ and collect terms, which results in

$$A_{nm} \ddot{q}_n + D_{nm} \dot{q}_n + B_{nm} q_n = 0$$

where

$$A_{nm} = \int_{\bar{x}}^1 \left[\frac{R^4 m}{m\Omega^2 R^4} \phi_n \phi_m \right] dx$$

$$D_{nm} = \frac{RD\Omega}{m\Omega^2 R^4} \phi_n' \phi_m' \Big|_{\bar{x}}$$

$$B_{nm} = \int_{\bar{x}}^1 \left[\frac{EI}{m\Omega^2 R^4} \phi_n'' \phi_m'' + \frac{R^2 G}{m\Omega^2 R^4} \phi_n' \phi_m' - \phi_n \phi_m \right] dx$$

Basis Functions For Assumed Mode Shapes

Polynomials are used as the basis functions, $\phi_n(x)$. Two sets of polynomials, meeting the necessary boundary conditions for hinged-free and cantilever-free beams, were used in this study. They are:

hinged-free:

$$\phi(x) = x$$

$$\phi(x) = x^6 - 2x^5 - \frac{5}{6}x^4 + \frac{10}{3}x^3 + x$$

cantilever-free:

$$\phi(x) = x^4 - 4x^3 + 6x^2$$

$$\phi(x) = x^5 - 10x^3 + 20x^2$$

Since these polynomials meet boundary conditions at $x=0$ and at $x=l$, and the blade formulation is integrated from $x = \bar{x}$ to $x=l$, the basis polynomials are transformed to new coordinates, where

$$x' = (1 - \bar{x})x + \bar{x}.$$

With this coordinate transformation, the new set of polynomials, which now fulfill the necessary boundary conditions at the hinge offset and at the blade tip, are now

hinged-free:

$$\phi(x) = x - \bar{x}$$

$$\begin{aligned} \phi(x) = & 1.48x^6 - 3.33x^5 - 0.12x^4 + 4.2x^3 - \\ & 0.79x^2 + 1.12x - 0.07 \end{aligned}$$

cantilever-free:

$$\phi(x) = 1.3x^4 - 5.2x^3 + 7.8x^2 - 0.92x + 0.03$$

$$\begin{aligned} \phi(x) = & 1.39x^5 - 0.44x^4 - 12.11x^3 + \\ & 25.10x^2 - 3.03x + 0.09 \end{aligned}$$

Appendix B. Exact Equations Of Motion For A Nonrotating Beam

The modal analysis assumes that the beam displacement is written as a sum of modal displacements:

$$Y(x, t) = R \sum_n \phi_n(x) q_n(t)$$

To find the exact analytic solution in the case of root constraint with both spring and damper, note that the boundary conditions are given by

$$\begin{aligned}\phi(0) &= 0 \\ \phi''(0) &= \left[\frac{KR}{EI} + \dot{\mu} \frac{DR}{EI} \right] \phi'(0) \\ \phi''(l) &= 0 \\ \phi'''(l) &= 0\end{aligned}$$

where K and D are spring and damper constants and all quantities are understood to refer to lagwise motion and are defined as in Appendix A.

These boundary conditions are satisfied by writing the mode summation equation as

$$\phi(x) = \phi_f(x) + \left[\frac{KR}{EI} + \dot{\mu} \frac{DR}{EI} \right] \frac{\phi_f'(0)}{\phi_c''(0)} \phi_c(x)$$

where $\phi_f(x)$ and $\phi_c(x)$ refer to hinged and cantilever mode shapes.

The hinged end mode shape solutions are given by

$$\begin{aligned}\phi_f(x) &= \cos(A) \sinh(Ax) + \cosh(A) \sin(Ax) \\ \phi_f(0) &= 0 \\ \phi_f'(0) &= A [\cos(A) + \cosh(A)] \\ \phi_f''(0) &= 0 \\ \phi_f''(l) &= A^2 [\cos(A) \sinh(A) - \cosh(A) \sin(A)] \\ \phi_f'''(l) &= A^3 [\cos(A) \cosh(A) - \cosh(A) \cos(A)]\end{aligned}$$

The cantilever mode shape solutions are given by

$$\begin{aligned}\phi_c(x) &= (\sin(A) - \sinh(A))(\sin(Ax) - \sinh(Ax)) + \\ &\quad (\cos(A) + \cosh(A))(\cos(Ax) - \cosh(Ax)) \\ \phi_c(0) &= 0 \\ \phi_c'(0) &= 0 \\ \phi_c''(0) &= -2A^2 [\cos(A) + \cosh(A)] \\ \phi_c''(l) &= -A^2 [1 + \cosh(A) \cos(A)] \\ \phi_c'''(l) &= A^3 [(\sin(A) - \sinh(A))(-\cos(A) - \cosh(A)) +\end{aligned}$$

$$(\cos(A) + \cosh(A))(\sin(A) - \sinh(A))]$$

Now use these known solutions for hinged and cantilever mode shapes in the combined solution given above:

$$\begin{aligned}\phi(0) &= 0 \\ \phi'(0) &= \phi_f'(0) \\ \phi''(0) &= [K + \dot{\mu}D] \left[-\frac{l}{2A} \right] \phi_c''(0) \\ \phi''(l) &= \phi_f''(l) + [K + \dot{\mu}D] \left[-\frac{l}{2A} \right] \phi_c''(l) \\ \phi'''(l) &= A^2 [\cos(A) \sinh(A) - \cosh(A) \sin(A) + \\ &\quad [K + \dot{\mu}D] \left[-\frac{l}{2A} \right] [-2A^2 [1 + \cosh(A) \cos(A)]]] \\ \phi'''(l) &= 0\end{aligned}$$

where

$$K = \frac{KR}{EI}$$

and

$$D = \frac{DR}{EI}$$

The boundary condition at the tip gives the eigenvalue equation:

$$\phi''(l) = 0$$

or

$$\begin{aligned}A^2 [\cos(A) \sinh(A) - \cosh(A) \sin(A)] + \\ A [K + \dot{\mu}D] [1 + \cosh(A) \cos(A)] = 0\end{aligned}$$

References

Ballin, M. G., and Dalang-Secretan, M. (1991). *Validation Of The Dynamic Response Of A Blade-Element UH-60 Simulation Model In Hovering Flight*. J. of the American Helicopter Society, Vol. 36, no. 4.

- Bramwell, A.R.S. (1976). Helicopter Dynamics. Edward Arnold Ltd., London, U.K.
- Curtiss, H.C. (1986). *Stability and Control Modeling*. Proceedings 12th European Rotorcraft Forum, Garmish Partenkirchen, Germany.
- Curtiss, H.C. and Zhao, X. (1988). *A Linearized Model of Helicopter Dynamics Including Correlation With Flight Test*. Paper presented at the 2nd International Conference on Rotorcraft Basic Research, College Park, MD.
- Diftler, M. A. (1988). *UH-60A Helicopter Stability Augmentation Study*. Paper presented at the 14th European Rotorcraft Forum, Milan, Italy.
- Goldberg, David E. (1989). Genetic Algorithms in Search, Optimization, and Machine Learning. Addison-Wesley Publishing Co., Reading, MA.
- Ham, N. D., Behal, B., and McKillip, R. M., Jr. (1983). *Lag Damping Augmentation Using Individual Blade Control*. Vertica, Vol. 7, No. 4.
- Holland, John H. (1975). Adaptation in Natural and Artificial Systems. The University of Michigan Press, Ann Arbor, Michigan.
- Johnson, W. (1980). Helicopter Theory. Princeton University Press, Princeton, New Jersey.
- Kaplita, T.K., Driscoll, J.T., Diftler, M.A., and Hong, S.W. (1989). *Helicopter Simulation Development By Correlation With Frequency Sweep Flight Test Data*. AHS 45th Annual National Forum, Boston, MA.
- Kim, F. D., Celi, R., and Tischler, M. B. (1990). *High-Order State Space Simulation Models Of Helicopter Flight Mechanics*. 16th European Rotorcraft Forum, Glasgow, Scotland.
- Ljung, L. System Identification: Theory for the User. (1987). Prentice-Hall, Inc., Englewood Cliffs, NJ.
- Maine, R. E. and Iliff, K. W. (1985). *Identification of Dynamic Systems - Theory and Formulation*. NASA Reference Publication 1138.
- Mayo, J. R., Occhiato, J. J., and Hong, S. W. (1990). *Helicopter Modeling Requirements For Full Mission Simulation And Handling Qualities*. AHS 47th Annual National Forum.
- Milne, G. W. (1986). *Identification of a Dynamic Model of a Helicopter from Flight Tests*. Ph.D. Thesis, Stanford University.
- Starer, David (1990). *Algorithms for Polynomial-Based Signal Processing*. Ph.D. Thesis, Center For System Science, Yale University.
- Tischler, Mark B. (1987). *Frequency-Response Identification of XV-15 Tilt-Rotor Aircraft Dynamics*. Ph.D. Thesis, Stanford University.

Session 4

Understanding Visual Cues

