Damage-mitigating control of mechanical structures: experimental verification of the concept

Sekhar Tangirala, Asok Ray† and Marc Carpino

The Pennsylvania State University, University Park, PA 16802, USA

Abstract. The concept of damage-mitigating control is built upon the two disciplines of *systems science* and *mechanics of materials*, and its goal is to achieve optimized trade-off between dynamic performance and structural durability of the plant under control. Simulation studies reported in recent publications show a substantial reduction of damage accumulation in the critical components of a rocket engine with no significant loss of performance. This paper reports experimental verification of the damage-mitigating control concept on a laboratory testbed which is a two-degree-of-freedom mechanical system excited by a computer-controlled shaker table. Test results demonstrate: (i) the important feature of optimized damage-mitigating control by extending fatigue life up to three and one half times with no significant performance degradation; and (ii) close agreement between the analytical prediction of damage and experimental observations.

1. Introduction

A major objective in control of complex mechanical systems such as electric power plants, advanced aircraft, and spacecraft is to achieve the mission objectives with increased reliability, availability, component durability, and maintainability. Therefore, performance, cost of maintenance and operation, and service life should be taken into consideration during the design process. In view of high performance requirements and availability of improved materials that may have significantly different damage characteristics relative to conventional materials, the lack of appropriate knowledge about the properties of these materials will lead to either of the following: (i) less than achievable performance due to overly conservative design; or (ii) over-straining of the mechanical structures leading to unexpected failures and drastic reduction of the service life. From this perspective, Ray et al (1994a) have proposed a methodology to achieve optimized trade-off between the system performance and structural durability of the plant under control. The concept of damage-mitigating control, which addresses the two disciplines of *systems science* and *mechanics of materials*, has been recently introduced in a two-part paper (Ray et al 1994b, c) along with simulation examples. Part I developed a continuous-time model of fatigue damage dynamics in the state-variable setting. This is in contrast to the usual notion of expressing the fatigue damage rate relative to the number of cycles. A unique advantage of this damage model

† Author to whom any correspondence should be sent; e-mail: a2r@ecl.psu.edu.

in the continuous-time setting is that it can be directly incorporated into the control system structure to provide the necessary information for on-line damage-mitigating control as well as for off-line synthesis of a control law. An optimal policy was formulated in Part II for open-loop control of a reusable rocket engine where the objective is to synthesize a robust control and diagnostic system by maintaining the damage rate and accumulation in the critical components within prescribed limits, while allowing for external disturbances and uncertainties in modeling of plant dynamics and damage dynamics. Simulation experiments showed a substantial reduction in the accumulated damage for an insignificant degradation in the rocket engine performance.

This paper reports experimental verification of the damage-mitigating control concept on a laboratory testbed which is a two-degree-of-freedom mechanical system excited by a computer-controlled shaker table. Test results demonstrate: (i) the important feature of optimized damage-mitigating control by extending fatigue life up to three and one half times with no significant performance degradation; and (ii) close agreement between the analytical prediction of damage and experimental observations.

2. The damage-mitigating control system

The concept of damage-mitigating control is most effective in enhancing the structural durability of a plant component when the component contains degrees of freedom which are not directly controlled by the plant inputs. An example is damage reduction in the turbine blades during

the upthrust transient motion of a rocket engine where the oxidant flow into a preburner is controlled to meet the engine performance requirements (Ray *et al* 1994a). Here, the torque of the turbopump, which is dependent on preburner gas temperature and pressure as well as on the main combustor pressure, is only indirectly affected by the oxidant flow. Since the turbine blade vibrations are not directly controlled, these blades are subjected to damage even during normal operation. In contrast to the conventional control system synthesis which only involves optimization of the system performance, the goal here is to achieve the specified performance of the engine while extending its service life.

The major tasks in the synthesis of a damage-mitigating control system are: (i) dynamic modeling of the mechanical system to be controlled; (ii) formulation of a damage model for each of the identified critical components; and (iii) synthesis of a control policy to achieve structural durability and high performance. The system and damage dynamics are represented in the continuous-time setting as follows:

Task period : starting time t_0 to final time t_f

System dynamics : $\dfrac{dx}{dt} = f(x(t), u(t))$ (1)

$x(t_0) = x_0$

Damage dynamics : $\dfrac{dv}{dt} = h(v(t), q(x, t))$ (2)

$v(t_0) = v_0 \; h \geq 0 \; \forall t \geq t_0$

where $x \in \mathbb{R}^n$ is the plant state vector; $u \in \mathbb{R}^m$ is the control input vector; $v \in \mathbb{R}^r$ is the damage state vector; and $q \in \mathbb{R}^p$ is the load vector yielded by the structural model. Equations (1) and (2) must satisfy the local Lipschitz condition (Vidyasagar 1992). The state variable representation of the system and damage dynamics in equations (1) and (2) facilitates formulation of the optimization problem. Given the initial conditions x_0 and v_0, the plant state propagation equations and the damage equations allow recursive computation of the state variables. For example, if the damage vector is used to represent a crack length, the initial damage state could include the initial crack size in the specimen under consideration. This initial damage state is important due to its significant effects on the dynamics of the non-linear damage model. The initial flaw size and material properties of any given specimen are inherently random, and inclusion of this randomness in equation (2) leads to stochastic differential equations. Reliability analysis and damage prognosis via solution of these stochastic differential equations are a subject of current research, and some of the results are reported in a forthcoming publication (Ray *et al* 1995).

3. Description of the laboratory testbed

A laboratory testbed has been constructed to verify experimentally the damage-mitigating control concept and to establish its feasibility. The testbed consists of a simple two-degree-of-freedom (DOF) mass–beam system excited by the vibratory motion of a shaker table as shown in figure 1. One DOF is directly controlled by the shaker
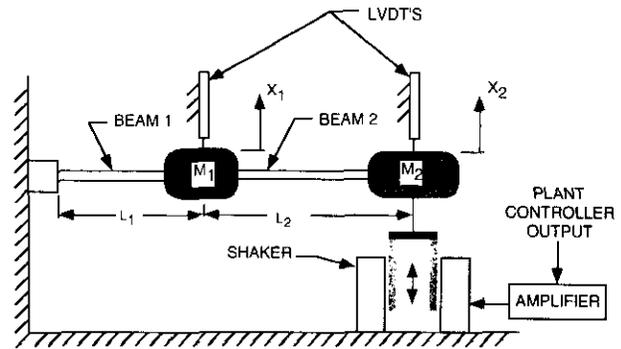

Figure 1. Mechanical structure of the laboratory testbed.

table whereas the other DOF is not controlled, but it is an observable DOF via position measurements of the vibrating masses. The test specimen, subjected to fatigue, is representative of mechanical damage in a plant component. The testbed is logically separated into two subsystems: (i) the plant subsystem consisting of the mechanical structure including the test specimen to be damaged, actuators, and sensors; and (ii) the control and communications subsystem consisting of computers, data acquisition and processing, and communications hardware. The simplicity of the construction of the testbed facilitates easy reconfiguration into more complex multi-input multi-output systems with additional degrees of freedom.

3.1. The plant subsystem

The discussion here focuses on the basic testbed configuration in figure 1 where the mechanical system consists of two discrete masses supported on cantilever beams. The second mass is connected to a shaker table which acts as the actuator for the system through a flexible pivot. The flexible pivot rigidly links the translating mass in the actuator to the mass on the beam in the direction of motion. The translational positions of the masses are measured using linear variable differential transformer (LVDT) sensors which are located at the midpoints of the masses M_1 and M_2. In the current configuration, the geometric parameters of the mechanical system are as follows:

Masses: $M_1 = 7.664$ lbm (0.238 slug); $M_2 = 2.702$ lbm (0.0839 slug).

Beam 1: length $L_1 = 11.84$ in; height $H_1 = 0.437$ in; width $W_1 = 0.870$ in; Young's modulus $E_1 = 30 \times 10^6$ psi.

Beam 2: length $L_2 = 8.5$ in; height $H_2 = 0.118$ in; width $W_2 = 0.437$ in; Young's modulus $E_1 = 30 \times 10^6$ psi.

3.2. The failure site

Since the objective of this testbed is to investigate the influence of different control policies on specific modes of fatigue failure in a dynamic setting, a structurally weakened element which is representative of a critical plant component is introduced to ensure the occurrence of an observable failure. In the two-mass configuration, two parallel failure sites are introduced by drilling a 0.312 inch hole located 2.66 inches from the center of mass M_1 as
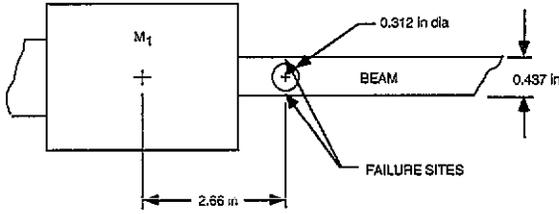
Figure 2. Top view of the failure site on the beam specimen.

shown in figure 2. As the specimen is subjected to cyclic stresses, cracks initiate on the upper and lower surfaces of the beam and propagate through the thickness towards the neutral axis.

3.3. Modeling of the mechanical structure

A structural model of the mechanical system described above is formulated using a cubic interpolation polynomial to approximate the lateral displacement of the beams and a linear approximation for the lateral displacement of the masses. The major assumptions in the model formulation include:

Lumped representation of the beam masses.

Beams subjected to pure bending (i.e., negligible torsional effects).

No deformation of the masses M_1 and M_2.

The assumption that the masses, M_1 and M_2, essentially behave as rigid bodies is justified by their relatively high stiffness. The actuator is assumed to force a point mass, M_2, at the tip of the beam 2 as seen in figure 1. The validity of the other assumption of negligible torsional effects is demonstrated later when the model is compared with actual experimental data.

The displacements of the beams 1 and 2, and the mass M_1 are approximated by the following relationships:

$$y_i(x) = C_{1i}\frac{x^3}{6} + C_{2i}\frac{x^2}{2} + C_{3i} + C_{4i} \qquad i = 1, 2$$

for the beams 1 and 2   (3a)

$$y_3(x) = C_5 x + C_6 \qquad \text{for the mass } M_1. \qquad (3b)$$

The ten constants in equations (3) are evaluated by applying standard boundary and compatibility conditions and using the physical dimensions listed above to yield the following dynamic equations of motion:

$$M_1\ddot{y}_1 + 244.77\bar{y}_1 - 18.38\bar{y}_2 = 0$$
$$M_2\ddot{y}_2 - 18.38\bar{y}_1 + 28.68\bar{y}_2 = F \qquad (4)$$

where $\bar{y}_1$ represents the lateral position of the center of mass, M_1; $\bar{y}_2$ represents the lateral position of the point mass, M_2; and F represents the force applied by the actuator on the point mass, M_2. The positions $\bar{y}_1$ and $\bar{y}_2$ correspond to the locations of the position sensors on the experimental testbed. An eigenvalue analysis of this system indicates that the natural undamped frequencies of the first and second modes of vibration are 9.73 Hz and 17.89 Hz,

respectively. The stresses at the failure sites on the beam 2 shown in figure 2 are given by

$$\sigma_{FS} = 5.47 \times 10^5 \bar{y}_1 - 2.58 \times 10^5 \bar{y}_2 \text{ psi.} \qquad (5)$$

Equation (5) accounts for the increased local stress at the critical point due to the reduction in the area moment of inertia caused by the hole drilled in the specimen as seen in figure 2. The model in equation (4) represents the dynamics of the mass–beam system where the damping terms are assumed to be insignificant. The stresses computed via equation (5) show excellent agreement with experimentally measured stresses. The displacement sensors are modeled as pure gains (i.e., having negligible lag). The shaker table and the signal amplifier unit are fairly non-linear, and the dynamics of the shaker are influenced by the loading effects of the mass–beam system. The problem of determining the dynamic relationship between the applied voltage and the force generated in the actuator system was circumvented by using system identification techniques. Time-domain frequency-sweep excitation (Ljung 1987) and experimentally measured data were used to identify a sixth-order voltage (actuator signal)-to-voltage (position sensor signal) transfer function. This data and the corresponding model outputs are shown in figure 3. This experimentally identified model is used to represent the plant dynamics of the system, and equation (5) models the time-dependent stress at the failure site.

3.4. Modeling of fatigue damage

The fatigue damage model in the continuous-time setting, reported by Ray and Wu (1994a), is first derived based on linear damage accumulation using a combination of Coffin–Manson and Basquin relationships (Bannantine et al 1990). Then, this linear damage model is modified following the damage curve approach (Bolotin 1989) to account for dependence of the damage rate on the current level of accumulated damage. Only the essential features of this fatigue damage model are presented in this section. Converting the strain amplitudes into stress amplitudes from the cyclic stress–strain curve, the rates of both elastic damage δ_e and plastic damage δ_p (Suresh 1991) are computed through differentiation as follows. If $\sigma \geq \sigma_r$, then

$$\frac{d\delta_e}{dt} = 2\frac{d}{d\sigma}\left(\left(\frac{\sigma - \sigma_r}{2(\sigma'_f - \sigma_m)}\right)^{-1/b}\right)\frac{d\sigma}{dt}$$

$$\frac{d\delta_p}{dt} = 2\frac{d}{d\sigma}\left(\frac{1}{\varepsilon'_f}\left(\frac{\sigma - \sigma_r}{2K'}\right)^{1/n'}\left(1 - \frac{\sigma_m}{\sigma'_f}\right)^{-c/b}\right)^{-1/c}\frac{d\sigma}{dt}. \qquad (6)$$

Otherwise, $d\delta_e/dt = 0$ and $d\delta_p/dt = 0$ where the current stress σ and the stress rate $d\sigma/dt$ are obtained from the structural model; σ_r is the reference stress obtained using the rainflow method (Rychlik 1993); $\sigma_m = (\sigma + \sigma_r)/2$ is the mean stress; and σ'_f, ε'_f, b, c, K', n' are fatigue parameters under cyclic operations (Bannantine et al 1990, Boller and Seeger 1987). Table 1 lists the fatigue parameters of 1018 steel which is the material of the specimen used in this experimental study. The damage rate $d\delta/dt$ is obtained as
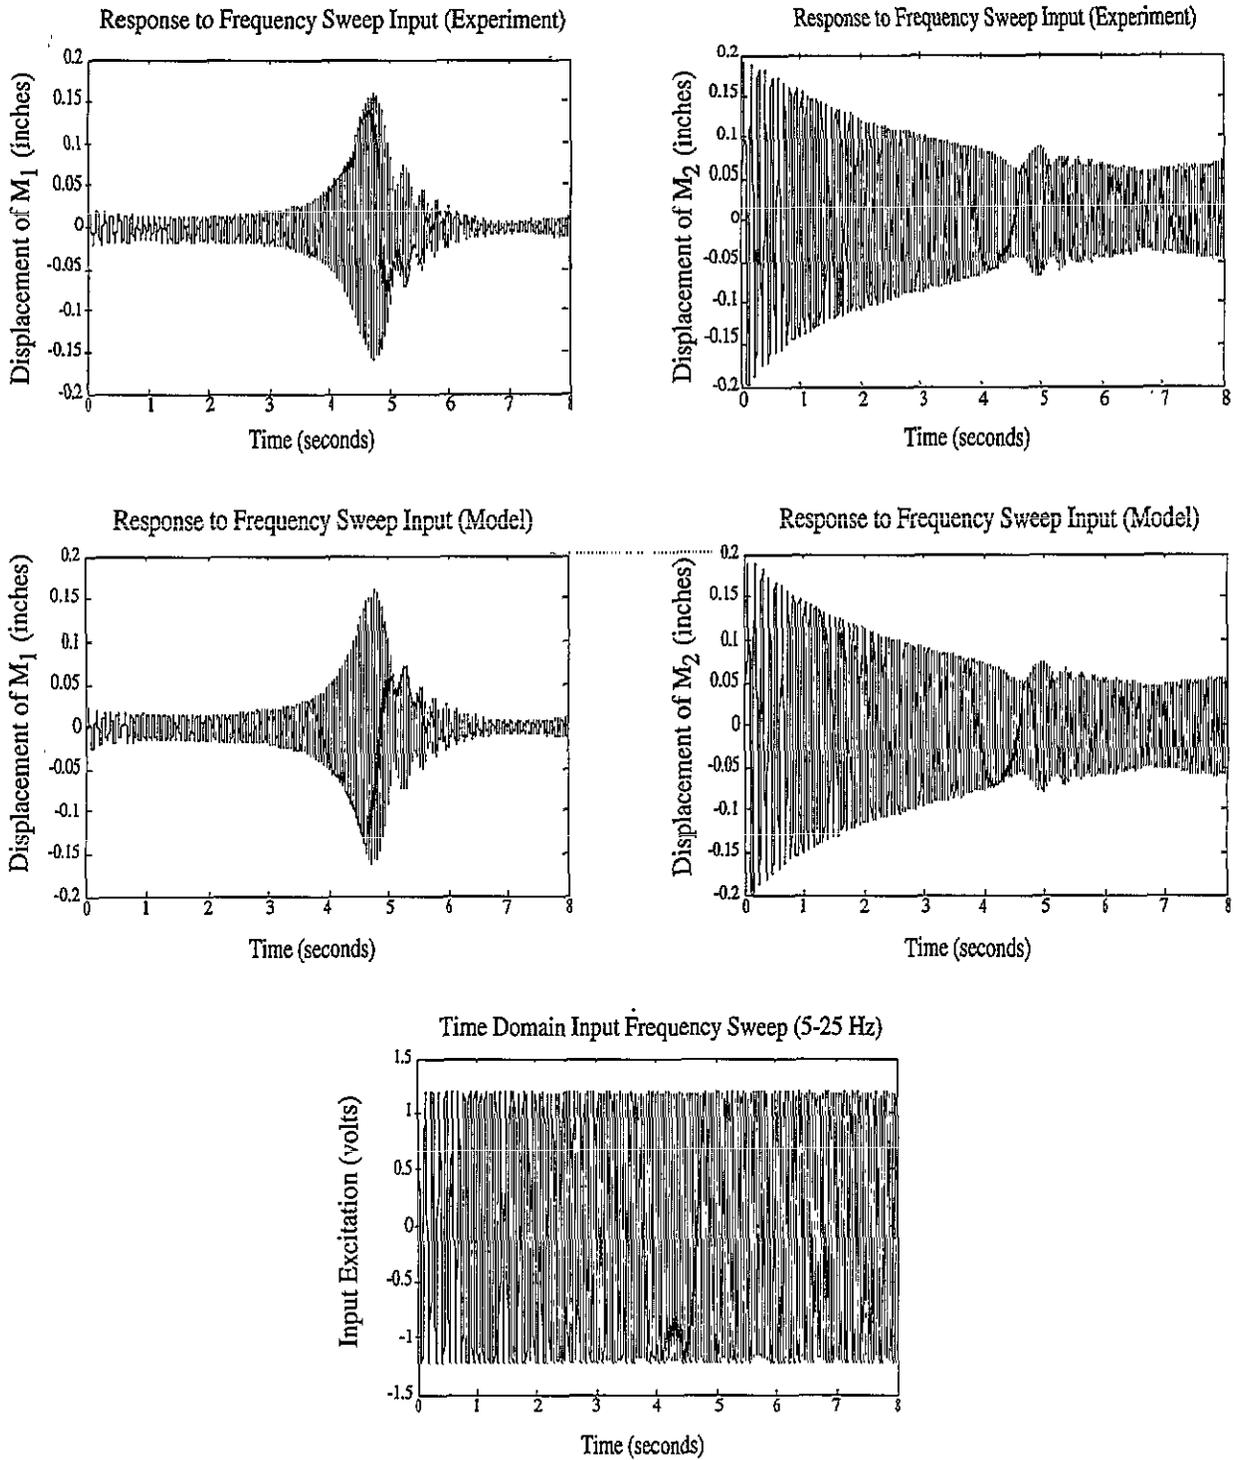
Figure 3. Comparison of experimental data and model response to frequency sweep input.

the weighted average of the elastic and plastic damage rates such that

$$\frac{d\delta}{dt} = w\frac{d\delta_e}{dt} + (1 - w)\frac{d\delta_p}{dt} \qquad (7)$$

where the weighting function, w, is selected as the ratio of the elastic strain amplitude and total strain amplitude.

Equations (6) and (7) are then used to obtain the damage rate at any instant. Since mechanical structures are generally subjected to loads of varying amplitude, equation (6) which is based on the linear rule of damage accumulation will lead to erroneous results due to the

sequence effect (Suresh 1991).Therefore, the linear damage is modified via a non-linear damage rule as follows

$$D = \delta^{\gamma(\sigma_a, D)} \qquad (8)$$

where D and δ are the current states of non-linear and linear damage accumulation, respectively, and σ_a is the stress amplitude. Equation (8) generates the following incremental change:

$$\Delta D = \gamma \left(D^{(\gamma-1)/\gamma}\right)\Delta\delta + \left(\frac{D(\ln D)}{\gamma}\right)\Delta\gamma. \qquad (9)$$

Table 1. Fatigue parameters for 1018 steel.

Strength coefficient, K'	194 ksi
Strain hardening exponent, n	0.226
Modulus of elasticity, E	30 000 ksi
Fatigue strength exponent, σ'_f	162 ksi
Fatigue ductility coefficient, ε'_f	0.338
Fatigue strength exponent, b	−0.11
Fatigue ductility exponent, c	−0.48

It follows from a crack propagation model such as the Paris model (Paris and Erdogan 1963) that the crack growth rate is dependent not only on the stress amplitude but also on the current crack length. Since the characteristics of γ in equation (9) may strongly depend on the type of the material, availability of pertinent experimental data for the correct material is essential for damage-mitigating control. An approach to evaluating γ at selected discrete levels of stress amplitude by interpolation based on the experimental data is given by Ray and Wu (1994a, b).

4. Open-loop control synthesis and simulation

The control policy is generated via non-linear programming which consists of optimizing a cost functional subject to constraints (Ray et al 1994a, c). These constraints are both physically motivated, such as actuator limits, and based on damage rate and accumulation in the critical components.

4.1. Cost functional

The non-linear optimization is realized via minimization of a cost functional that includes performance specifications subject to physical and damage constraints. The optimization yields an optimal control sequence, $\{u_k\}$, that manoeuvres the plant from a known initial condition to a specified final state, without violating any of the constraints. The quadratic cost functional is chosen to be the weighted ℓ_2 norm of the error in selected plant outputs, control inputs and the final output errors. The cost functional is of the form

$$
\begin{aligned}
J(u_k) &= [y_N - y_{ss}]^T Q_f [y_N - y_{ss}] \\
&+ \sum_{k=0}^{N-1} \left([y_k - y_k^{nom}]^T Q [y_k - y_k^{nom}] \right. \\
&\left. + [u_k - u_k^{nom}]^T R [u_k - u_k^{nom}] \right)
\end{aligned}
\tag{10}
$$

where N is the total number of discretized steps which represent the period from the initial time, t_0, to the final time, t_f; $\{y_k^{nom}\}$ is the desired nominal trajectory of the performance variable y; y_{ss} is the steady-state value of $\{y_k^{nom}\}$; the penalty matrices, Q, Q_f and R are symmetric positive definite having compatible dimensions. The minimization is carried out subject to constraints. These constraints take the form of actuator limits and limits on the damage rate experienced by the material at the failure site during operation. The damage constraints have not been applied to the accumulated damage, v_k, since the actual damage accrued in any given cycle is very small and

varies with specimen fatigue. Reducing the damage rate in a cycle of operation results in the reduction of the damage accumulated during that cycle, leading to an increase in life of the system. Due to the optimization, this increase in life can be achieved with only a marginal loss of performance. The weight on the control inputs is included to avoid large fluctuations of the input signal as this could lead to chatter and increased fatigue damage. The relative weights on the two components of the cost functional are selected by the design engineer to emphasis performance, transient plant behavior, damage dynamics or control effort as required. It should be noted, however, that increasing the damage rate constraint has a direct impact on the system performance. So, a judicious selection of the constraints and weights can result in, within limits, different levels of damage mitigation with the associated loss of performance.

If the plant model is completely controllable, then there exists a control sequence that can manoeuvre the plant from any arbitrary initial state to any specified final state. However, this does not guarantee that a solution exists for every constrained optimization problem. Whenever the admissible control inputs are restricted to a feasible set, certain final states may not be accessible without violating some of the constraints. In this work, a general purpose non-linear programming software package NPSOL (Gill et al 1991) has been used to solve the open-loop constrained optimization problem. The details are reported by Ray and Wu (1994b) and Ray et al (1994a).

4.2. System constraints

Constraints on the control input sequence, $\{u_k\}$, represent actuator saturation and hard limits on the actuator stroke. Additionally pre-assigned limits are imposed on the damage rate to assure structural durability. These conditions take the following form:

$$
\text{Natural bounds}: 0 \leq u(t) \leq \alpha \tag{11}
$$

$$
\text{Damage constraints}: v(t) \leq \beta(t) \text{ for } t \in [t_0, t_f] \tag{12}
$$

where t_0 is the initial time and t_f the final time; $\alpha \in \mathbb{R}^m$ is the actuator limit and $\beta \in \mathbb{R}^r$ is the specified tolerance for damage rate.

4.3. Problem formulation

The problem is to generate, for the mass–beam structure shown in figure 1, an optimal control law that will make a trade-off between the conflicting requirements of performance and damage mitigation over the time interval $[t_0, t_f]$ This optimization problem is represented in the discrete time setting to find an optimal control sequence $\{u_k\}$ that will manoeuvre the plant from a known initial condition to a prescribed final state subject to the following constraints:

(i) The discretized version of the continuous-time dynamic system in equation (4);

(ii) The actuator limits in equation (11); and
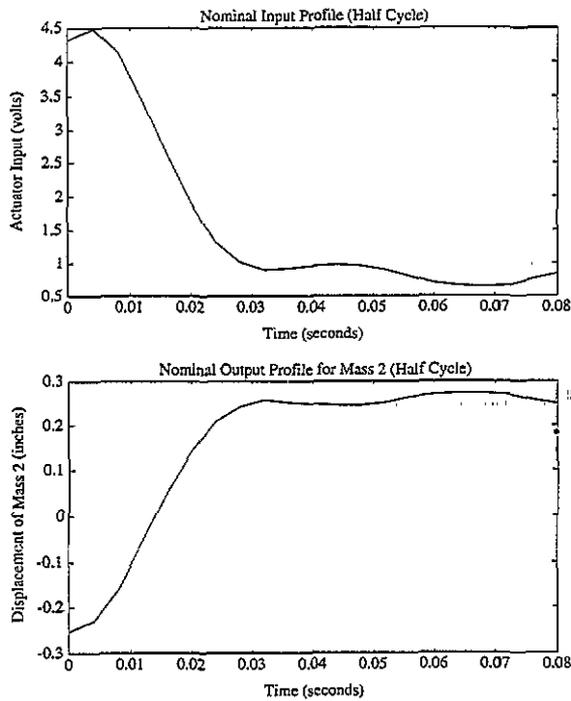
(iii) The constraints on the damage rate in equation (12).

Figure 4. Nominal input and output profiles for the optimization.

Figure 5. Constrained actuator input.

Figure 6. Comparison of constrained and unconstrained outputs (simulation).

In the mechanical system of figure 1, the voltage input to the shaker table directly controls the position of the mass M_2. The frequency of operation of the shaker table is 5.95 Hz (i.e., the full-cycle period of approximately 168 milliseconds), which is one third of the second natural frequency of the system. Note that the third harmonic of this frequency excites resonance of the mass M_1 which is not directly controlled by the shaker table. The optimization is carried out over one half cycle of motion with the displacement of M_2 chosen as the performance variable. The nominal output profile is determined by the closest approximation to a half square wave that the displacement trajectory of M_2 can achieve on the actual dynamic system over one half cycle of operation. This profile and the corresponding nominal excitation are shown in figure 4. Resonance of M_1 causes fatigue damage of the test specimen shown in figure 2. This is a direct consequence of the high local stresses at the failure site caused by the large displacements of M_1. This configuration is, therefore, representative of a real physical situation where the specification of performance requirements in one subsystem of a dynamic system may lead to fatigue damage in other subsystems.

5. Simulation and experimental results

The optimization task described in the previous section was carried out over one half cycle of motion of the mass–beam structure. The sampling rate was selected to be 4 ms, which is achievable on the 486-based control computer of the experimental facility. This sampling rate generates a total of 21 discrete time steps to manoeuvre the mass M_2 from an initial position of -0.25 inches to a final position of $+0.25$ inches over one half of the period of operation. Simulation experiments show that the maximum damage
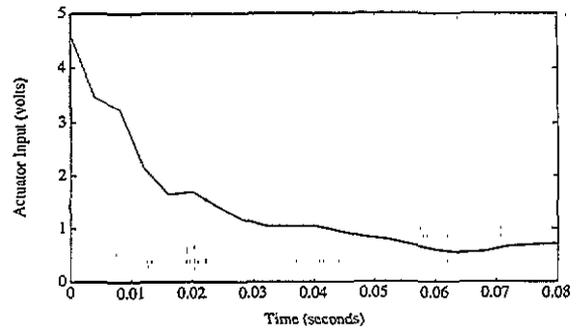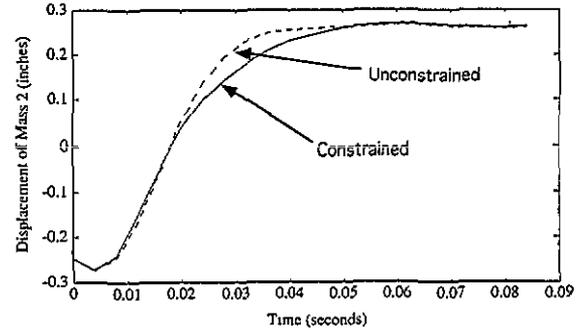
rate experienced at the failure site due to this manoeuvre is of the order of 1.0 s^{-1}. A damage rate constraint of 0.05 s^{-1} is therefore chosen for the optimization task. It should be noted, however, that the initial damage of the specimen is a random variable (Ray *et al* 1995) which depends on several factors like the machining quality, the heat treatment, size of the notch, etc. The level of initial damage used in the optimization has a significant effect on the damage rate and accumulation due to non-linearities in the damage dynamics. The expected value of the initial damage has been assumed to be 0.0005 in the simulation, and the resulting life predictions match experimental data.

The γ-parameters in equations (8) and (9) at different levels of damage need to be obtained for various stress amplitudes based on experimental data of crack growth and linear damage prediction data generated by equations (6) and (7). As the experimental data of short crack growth were not available for 1018 steel which is the material of the test specimen, the γ-parameters in the non-linear damage model of 1018 steel were evaluated in the appropriate stress range after adjusting those for 4340 steel that were reported by Ray and Wu (1994a, b). Although the predicted damage may not be accurate in the absolute sense due to possibly inexact γ-parameters, the profile of damage dynamics generated by this model shows the right trend. Therefore, this damage information is considered to be sufficient for comparison of the results of unconstrained and constrained optimization.

The constrained optimization yields the input waveform shown in figure 5. The corresponding output profile is plotted along with the unconstrained (nominal) output in figure 6. This plot shows that the loss of performance, as defined by the deviation from the nominal profile, is quite small. Figure 7 shows the fast Fourier transform (FFT) of
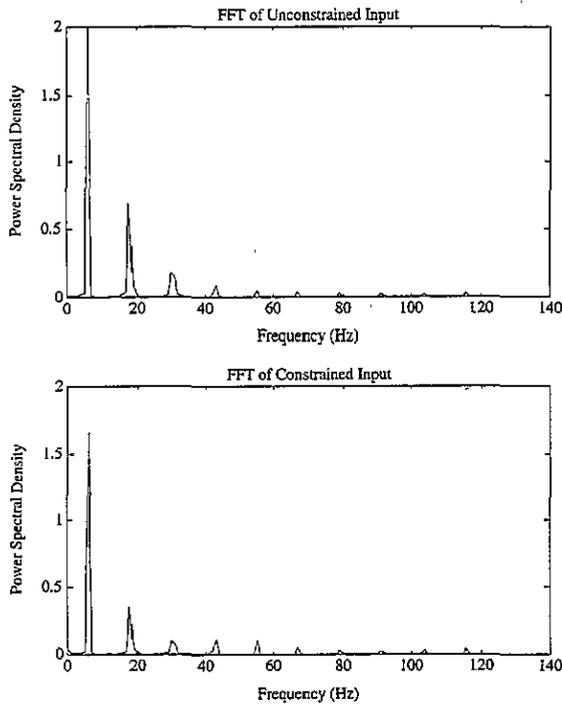
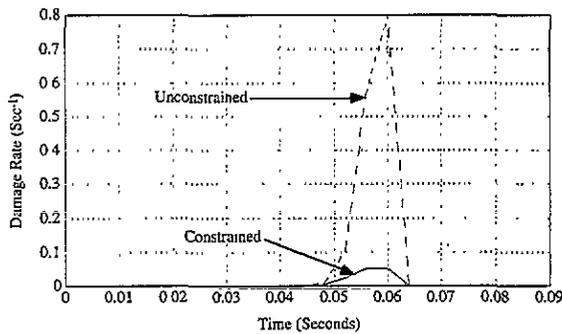**Figure 7.** FFTs of constrained and unconstrained inputs.

Figure 8. Comparison of constrained and unconstrained damage rates.

Figure 9. Comparison of simulated and experimental positions of the mass M_2.

the constrained and unconstrained input waveforms. The FFT of the constrained waveform shows a marked reduction in the magnitude of the third-harmonic component at ~18 Hz relative to the unconstrained waveform. As mentioned earlier, this third harmonic excites resonance of the mass M_1, and is the major factor causing fatigue damage at the failure site. Simply removing this component from the input excitation, however, leads to a considerable loss of performance which is unacceptable. Therefore, the magnitude of this third-harmonic component is substantially reduced while some of the other (apparently harmless) harmonics are moderately increased to minimize the loss of performance. Figure 8 compares the damage rates obtained from the simulation of the dynamic system using the unconstrained and constrained input waveforms. A considerable reduction in the damage rate due to the constraint is clearly observed. The next two paragraphs present simulation and experimental results to demonstrate how the constrained optimization procedure achieves a substantial increase in the specimen life with only a marginal loss of performance.

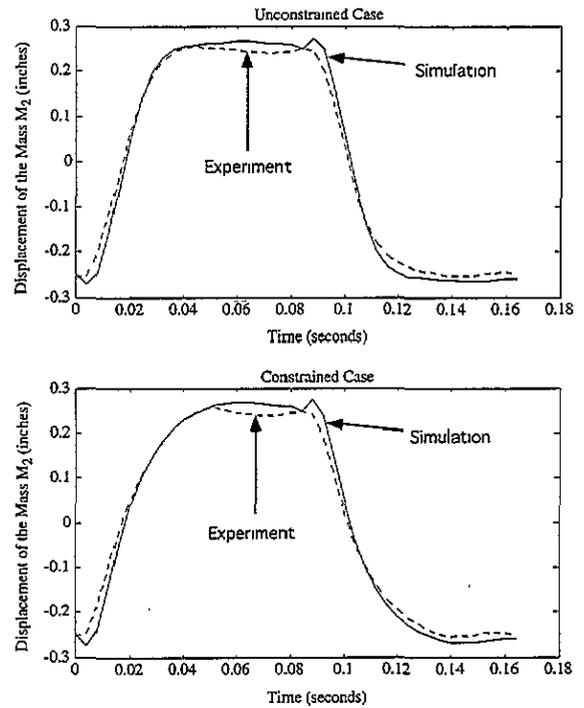Simulation of the dynamic system to the point of

failure of the specimen (i.e., when the damage accumulation reaches unity) yields a predicted life of 3639 cycles for the unconstrained case and 11 814 cycles for the constrained case. These simulation results show an increase in life by a factor of ~3.25 while the loss of performance is insignificant as seen in figure 6. It should be noted, however, that the predicted life is based on an expected value of 0.0005 for the initial damage and therefore, this approach may at most represent the specimen life in a statistical sense (Ray *et al* 1995).

The constrained and unconstrained input sequences, generated earlier for the simulation, were also used to excite the shaker table and mechanical structure in the testbed. These experiments were conducted using identical specimens that were machined from the same stock of 1018 steel. Figure 9 compares the output signals (i.e., the displacement of the mass M_2) obtained from the simulation and experimental results when excited by the unconstrained and constrained input signals. Table 2 summarizes the results of several experiments with both unconstrained and constrained inputs. The mean value of specimen life is 3767 cycles for the unconstrained case and 12 434 for the constrained case. These results show an increase in life by a factor of ~3.30 which agrees very closely with the corresponding gain of ~3.25 predicted by simulation. Table 2 shows that there is some scatter in the number of cycles to failure for different specimens. This is due to the stochastic nature of the damage process and the randomness in the initial flaw size in apparently identical specimens.

6. Summary and conclusions

This paper presents an extension of the work reported in a recent two-part paper which introduced the concept of

Table 2. Summary of experimental results for constrained and unconstrained cases.

Specimen number	Cycles to failure
1	3934 (unconstrained)
2	3812 (unconstrained)
3	3554 (unconstrained)
4	11 222 (constrained)
5	13 960 (constrained)
6	12 120 (constrained)

damage-mitigating control and demonstrated its feasibility by simulation of the upthrust transients of a rocket engine. A testbed has been constructed to verify this concept of damage mitigation experimentally. The experimental results on the testbed are in close agreement with the simulation results, and demonstrate an increase in the life of test specimens by a factor of about 3.5 with no significant loss of performance. The experimental data show evidence of random fatigue (Ray *et al* 1995). This suggests restructuring the governing equations for damage dynamics in the stochastic setting, and formulation of the control law via stochastic optimization.

Acknowledgments

References

Bannantine J A, Comer J J and Handrock J L 1990 *Fundamentals of Metal Fatigue Analysis* (Englewood Cliffs, NJ: Prentice-Hall)

Boller C and Seeger T 1987 *Materials Data for Cyclic Loading* (Amsterdam: Elsevier)

Bolotin V V 1989 *Prediction of Service Life for Machines and Structures* (Warrendale, PA: ASME Press)

Gill P E, Murray W, Saunders M A and Wright M H 1991 *NPSOL Version 4 06* Software Distribution Center, Stanford University, Palo Alto, CA

Ljung L 1987 *System Identification: Theory for the User* (Englewood Cliffs, NJ: Prentice-Hall)

Paris P C and Erdogan F 1963 A critical analysis of crack propagation laws *ASME J. Basic Eng.* **D85** 528–34

Ray A, Dai X, Wu M-K, Carpino M and Lorenzo C F 1994a Damage-mitigating control of a reusable rocket engine *AIAA J. Propulsion Power* **10** 225–34

Ray A, Tangirala S and Newman J C 1995 Stochastic modeling of damage dynamics for failure prognostics and risk analysis *American Control Conf. (Seattle, WA)*

Ray A and Wu M-K 1994a Fatigue damage control of mechanical systems *Smart Mater. Struct.* **3** 47–58

——1994b Damage-mitigating control of space propulsion systems for structural durability and high performance *NASA Lewis Research Center Contractor Report* 194470, March

Ray A, Wu M-K, Carpino M and Lorenzo C F 1994b Damage-mitigating control of mechanical systems: part I–conceptual development and model formulation *ASME J. Dynamic Syst., Meas. Control* **116** 437–47

——1993c Damage-mitigating control of mechanical systems: part II–formulation of an optimal policy and simulation *ASME J. Dynamic Syst., Meas. Control* **116** 448–55

Rychlik I 1993 Note on cycle counts in irregular loads *Fatigue Fract. Eng. Mater. Struct.* **16** 377–90

Suresh S 1991 *Fatigue of Materials* (Cambridge: Cambridge University Press)

Vidyasagar M 1992 *Nonlinear Systems Analysis* 2nd edn (Englewood Cliffs, NJ: Prentice-Hall)