JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 20, No. 1, January–February 1997

Flight Control Applications of ℓ_1 Optimization

Mark Spillman* and D. Brett Ridgely†
Air Force Institute of Technology, Wright–Patterson Air Force Base, Ohio 45433-7531

The ℓ_1 optimization method is used to handle tracking issues in the design of discrete flight controllers for a single-input/single-output aircraft longitudinal control problem. The ℓ_1 optimization approach is discussed theoretically and on a broad conceptual level. Constraints are developed to handle control deflection and rate limitations, overshoot and undershoot limitations, and steady-state error requirements. The effect of each constraint is evaluated with simulations of the aircraft model. A closed-loop model-matching design is also presented, which produces acceptable tracking results with a lower-order controller.

Introduction

A GREAT deal of research in optimal flight control design in recent years has focused on H_2 and H_∞ optimization.[1] Although both methods work well for specific classes of inputs, neither method adequately handles hard magnitude and time domain constraints on the system, such as control deflection limitations, control rate limitations, and overshoot restrictions in the system response. The ℓ_1 optimization method,[2] which minimizes the maximum magnitude of a system's output to an unknown but bounded magnitude input, can be used to incorporate hard magnitude constraints on the system. Because this optimization method is a time domain approach, it can also address time domain constraints on the system response.

Dahleh and Diaz-Bobillo[2] have done the most comprehensive work on ℓ_1 optimization to date. In their work, they pose the ℓ_1 optimization problem as a linear programming problem and solve it exactly for one-block systems. Three methods for finding approximate solutions to multiblock problems are also presented. Dahleh and Diaz-Bobillo propose several methods of incorporating control deflection limitations, control rate limitations, and overshoot; however, some implementation details are omitted, and few comparisons between the different methods are shown. The purpose of this paper is to investigate and compare different magnitude and time domain constraints that can be added to the ℓ_1 optimization problem to produce systems with good tracking characteristics. A one-block model matching design is also investigated to see whether a system with good tracking qualities can be obtained with a lower-order controller. MATLAB® software was used to conduct this research.

ℓ_1 Optimization

The ℓ_1 optimization problem was first introduced by Vidyasagar,[3] but Dahleh and Pearson[4] are responsible for its more general solution. The goal of this section is to explain Dahleh and Pearson's method of solution, which involves posing the problem as a linear programming problem. To simplify the explanation, the introductory development considers the case of one-block problems only. The changes necessary to find solutions to multiblock problems are discussed thereafter.

The discrete system in Fig. 1, where $r(k) \in \mathbb{R}^n$ is an exogenous input sequence of unknown but bounded magnitude and $m(k) \in \mathbb{R}^p$ is the output sequence to be controlled, represents the standard ℓ_1 problem. If $\Phi \equiv F_l(P, K)$ is the closed-loop transfer function from m to r, then the objective of ℓ_1 optimization can be written as

$$\inf_{K \text{ stabilizing}} \|\Phi\|_1 = \inf_{K \text{ stabilizing}} \left[\max_{1 \le i \le p} \sum_{j=1}^{n} \sum_{k=0}^{\infty} |\phi_{ij}(k)| \right] \tag{1}$$

Several steps must be taken to pose this as a linear programming problem. First, the nonlinear absolute value function in the norm calculation must be removed. This is accomplished by a standard change of variables used in linear programming. Let $\Phi = \Phi^+ - \Phi^-$, where Φ^+ and Φ^- are sequences of $p \times n$ matrices with positive entries. The norm calculation can then be replaced by

$$\max_{1 \le i \le p} \sum_{j=1}^{n} \sum_{k=0}^{\infty} \left[\phi_{ij}^+(k) + \phi_{ij}^-(k)\right] \tag{2}$$

which is equal to the norm if, for every (i, j, k), either ϕ^+ or ϕ^- is zero. Because ϕ^+ or ϕ^- must be zero at the optimal solution, this substitution is valid.

Before searching for the variables Φ^+ or Φ^- that minimize the one norm of Φ, constraints must be imposed to ensure that the resulting Φ will be stable and realizable. These two problems are handled with the Youla et al.[5] parameterization. Using this parameterization, Φ can be expressed as $\Phi = H - UQV$, where H, U, and V are known and Q is unknown (see Ref. 6 for expressions for H, U, and V). For one-block problems U and V are inner[6] and, thus, invertible. This means that Q can be solved for directly, $Q = U^{-1}(H - \Phi)V^{-1}$, which makes it easy to see that Q will be stable if and only if the transfer function $(H - \Phi)$ cancels the unstable zeros of U and V. In other words, if the unstable zeros of U and V are denoted as a_i, then Q will be stable if and only if $\Phi(a_i) = H(a_i)$, for $1 \le i \le N$. Further, if Φ is written as a function of λ, where $\lambda = z^{-1}$,

$$\hat{\Phi}(\lambda) = \sum_{k=0}^{\infty} \Phi(k)\lambda^k \tag{3}$$

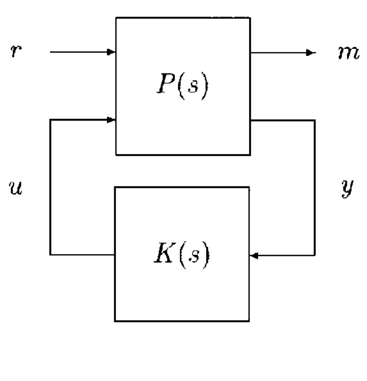

Fig. 1 Standard ℓ_1 optimization problem.

Received Aug. 30, 1995; revision received June 30, 1996; accepted for publication July 26,1996.

*M.S. Student, Aeronautics and Astronautics Department; currently Stability and Control Engineer, Wright Laboratory, WL/FIGC, Building 146, 2210 Eighth Street, Suite 21. Member AIAA.

†Associate Professor, Aeronautics and Astronautics Department, 2950 P Street, Building 640. Senior Member AIAA.

then this constraint (for the case of simple zeros) can be expressed as

$$\begin{bmatrix} 1 & a_1 & a_1^2 & \cdots \\ 1 & a_2 & a_2^2 & \cdots \\ 1 & \vdots & \vdots & \cdots \\ 1 & a_N & a_N^2 & \cdots \end{bmatrix} \begin{bmatrix} \hat{\Phi}(0) \\ \hat{\Phi}(1) \\ \hat{\Phi}(2) \\ \vdots \end{bmatrix} = \begin{bmatrix} H(a_1) \\ H(a_2) \\ \vdots \\ H(a_N) \end{bmatrix} \tag{4}$$

or $A_{\text{feas}}\Phi = b_{\text{feas}}$, which is linear in Φ.

With the specified modifications, the ℓ_1 optimization problem becomes

$$\inf \sum_{k=0}^{\infty} \hat{\Phi}^{+}(k) + \hat{\Phi}^{-}(k)$$

$$\text{subject to } \; A_{\text{feas}}[\hat{\Phi}^{+}(k) - \hat{\Phi}^{-}(k)] = b_{\text{feas}} \tag{5}$$

$$\hat{\Phi}^{+}(k) \geq 0, \qquad \hat{\Phi}^{-}(k) \geq 0$$

which is a linear programming problem with an infinite number of variables and a finite number of constraints. The corresponding dual problem has a finite number of variables and an infinite number of constraints. However, if there are no unstable zeros of U and V on the unit circle, at some large enough k, the a_i^k will be small enough that only a finite number of these constraints will be active.[2] Thus, the dual problem is finite dimensional, and an exact solution can be found. Further, the existence of a solution to the dual problem guarantees the existence of the same solution to the primal problem. In fact, the ℓ_1 optimization problem can be solved directly in the primal space by truncating the series in Eq. (5) at a large enough value.

There are a few modifications that must be made to the preceding formulation for multiblock problems. First, U and V may not be invertible. However, if the problem is nonsingular, i.e., all of the controls are penalized and no measurements are perfect, then a left inverse of U and a right inverse of V will exist, which is all that is required. Additionally, for multiblock problems, it is the left unstable zeros of U and the right unstable zeros of V that must cancel with zeros of $(H - \Phi)$.

Multiblock problems have an infinite number of variables as well as an infinite number of constraints and thus cannot be solved exactly. To counter this problem, Dahleh and Diaz-Bobillo[2] proposed three ways to find approximate solutions. The first method, known as the finitely many variables approach, constrains the polynomial solution Φ to a fixed length. The resulting compensator provides a suboptimal but feasible solution to the problem. The second method, known as the finitely many equations approach, truncates the number of dual variables, which is the same as solving the primal problem with a finite number of constraints. The solution to this problem is superoptimal and infeasible. The final and most viable method is known as the delay augmentation (DA) approximation. This method is generally considered the best method to use for multiblock problems since it carries more information about the optimal solution than the other two approaches.

The DA approach embeds the multiblock problem into a larger one-block problem by augmenting pure delays to U and V. The resulting one-block problem, which contains extra degrees of freedom in Q, can then be solved exactly. Whereas the solution to this problem is superoptimal and infeasible, it serves as an upper bound to the true optimal. To get a feasible solution, the extra degrees of freedom are simply stripped out of Q. The resulting solution is suboptimal but provides a lower bound to the optimal solution. Thus, this method produces both a feasible solution and bounds on the optimal solution.

In later sections, the standard ℓ_1 linear programming problem is augmented with constraints on the maximum magnitude of the controlled output to an exogenous step input. These additional constraints are posed using the vector ℓ_∞ norm,

$$\|m\|_\infty = \sup_k |m(k)| \tag{6}$$

Understanding ℓ_1 Optimization

To answer the question of how best to use ℓ_1 optimization, it is important to first understand what ℓ_1 optimization is trying to accomplish. The formal mathematical definition given in the preceding section is not important here; rather, a simple conceptual idea of how the method works is sufficient. By definition, ℓ_1 optimization attempts to minimize the absolute sum of a system's sampled pulse response. Conceptually, this optimization method works by pushing down on the pulse response from all sides. In other words, both peak-to-peak gains and long pulse responses are penalized because both tend to increase the absolute sum.

Because the primary interest is how best to use ℓ_1 optimization for tracking problems, it is instructive to examine the unit pulse and step responses of a simple discrete system. Consider the continuous system

$$H(s) = \frac{1}{s+1} = \left[\begin{array}{c|c} -1 & 1 \\ \hline 1 & 0 \end{array}\right] \tag{7}$$

which, discretized at $\frac{1}{3}$ Hz using a zero-order hold (ZOH), is equivalent to

$$H(z) = \left[\begin{array}{c|c} 0.05 & 0.95 \\ \hline 1 & 0 \end{array}\right] \tag{8}$$

The sampled unit pulse response of this system is shown in Fig. 2. The one norm of this system can be calculated by inspection. The total sum of samples 1–4 in Fig. 2 appears to be approximately 1. Indeed, the one norm for this system is 1. The unit step response of the system in Eq. (8) is also shown in Fig. 2. Notice the distinct relationship between the unit pulse response and the unit step response.

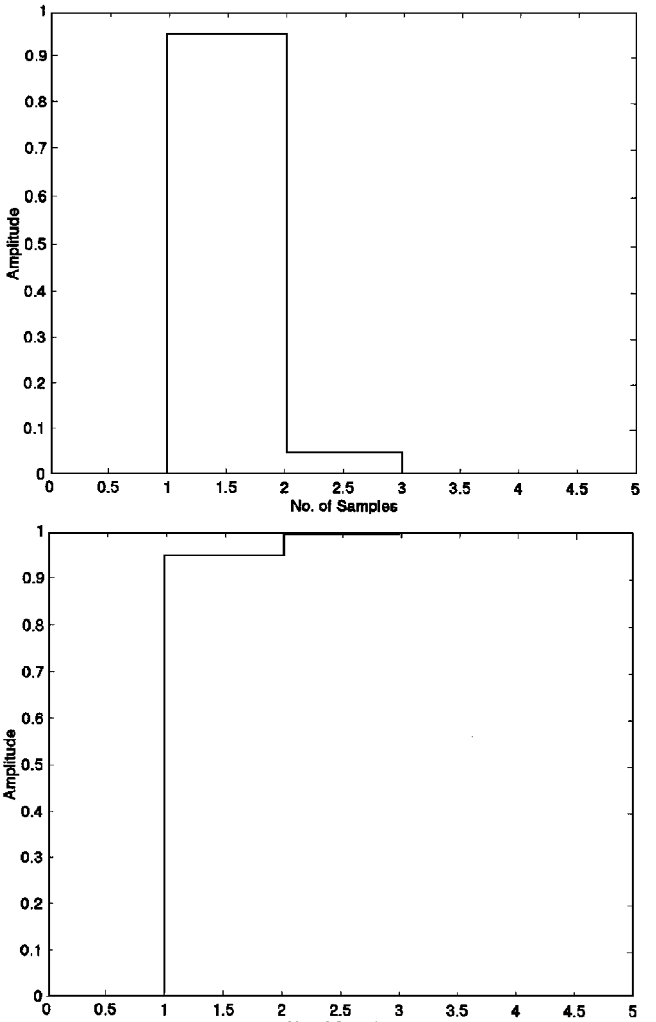

Fig. 2 Pulse (top) and step (bottom) response of a discrete system.

If r is the sampled unit step response and h is the sampled unit pulse response, then

$$r(k) = \sum_{j=1}^{k} h(j) \qquad \text{for} \qquad k = 1, 2, \ldots \tag{9}$$

This relationship implies that the faster the pulse response decays to zero, the quicker the step response reaches its steady-state value. Because ℓ_1 optimization penalizes long pulse responses, it logically also penalizes slow unit step responses. This fact and the general relationship between the unit pulse and unit step responses are particularly important in using ℓ_1 optimization for tracking problems.

Design Problem

Hereafter a single-input/single-output (SISO) longitudinal model of the AFTI F-16 is used to illustrate various tracking design issues. This particular design example has a large frequency spread in the system dynamics (fast actuator dynamics and slow phugoid modes), which makes for a rather challenging ℓ_1 optimization problem. The fast dynamics in the system force a fast discretization period to be chosen, whereas the slow dynamics ensure the system's pulse response will decay very slowly. These two factors lead to a large number of variables required to accurately describe the system's pulse response.

The tracking problem is to accurately command a 1-g (always from trim) normal acceleration of the aircraft. The stabilator is the only control surface considered in the model, and it is limited to ±25-deg deflection angle and ±60-deg/s deflection rate. A linear model of the aircraft is given in the Appendix.

All simulations are done with sampled-data systems, i.e., discrete controllers with continuous system models. Step inputs (from trim) of 1-g normal acceleration, applied 1 s after simulations are started, are used to evaluate tracking performance.

Sensitivity Minimization

The goal of most tracking problems is to minimize the error between the commanded input and the system output. This type of problem can be posed as a sensitivity minimization problem, such as the one depicted in Fig. 3. For the AFTI F-16 problem, the exogenous input r is an unknown commanded normal acceleration input with maximum magnitude less than or equal to one, and the controlled output m is the weighted error between the commanded acceleration and the actual aircraft acceleration. K is the unknown compensator, and G is the unweighted plant, described in the last section.

W_s is a weighting, which can be used to minimize the error to a select frequency range of command signals. Choosing an appropriate W_s, however, is often a difficult task. For this reason, W_s is set to 1 and alternative methods are explored to replace the frequency weighting.

An ℓ_1 optimization was performed on the system in Fig. 3. Because a sensitivity problem places no penalty on control usage, a small penalty, $\mu = 1 \times 10^{-5}$, was added to ensure the left inverse of U exists. The optimal closed-loop system has a one norm of 2.07 and the controller is fourth-order. The system response to a 1-g step in normal acceleration is shown in Fig. 4. The jags in the response are a product of the sampled-data simulation. As the sample rate increases the jags become less apparent.

Notice that the step response is extremely fast. This is mainly because there was only a small penalty placed on control usage. However, as discussed in the preceding section, unconstrained ℓ_1 optimization tends to produce very quick step responses. Plots of

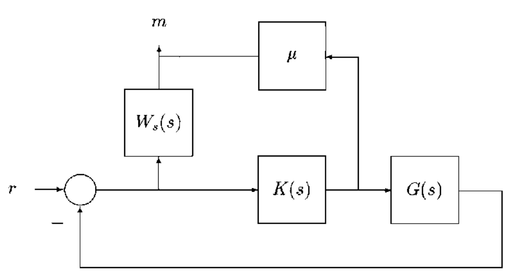

Fig. 3 Block diagram of ℓ_1 sensitivity.

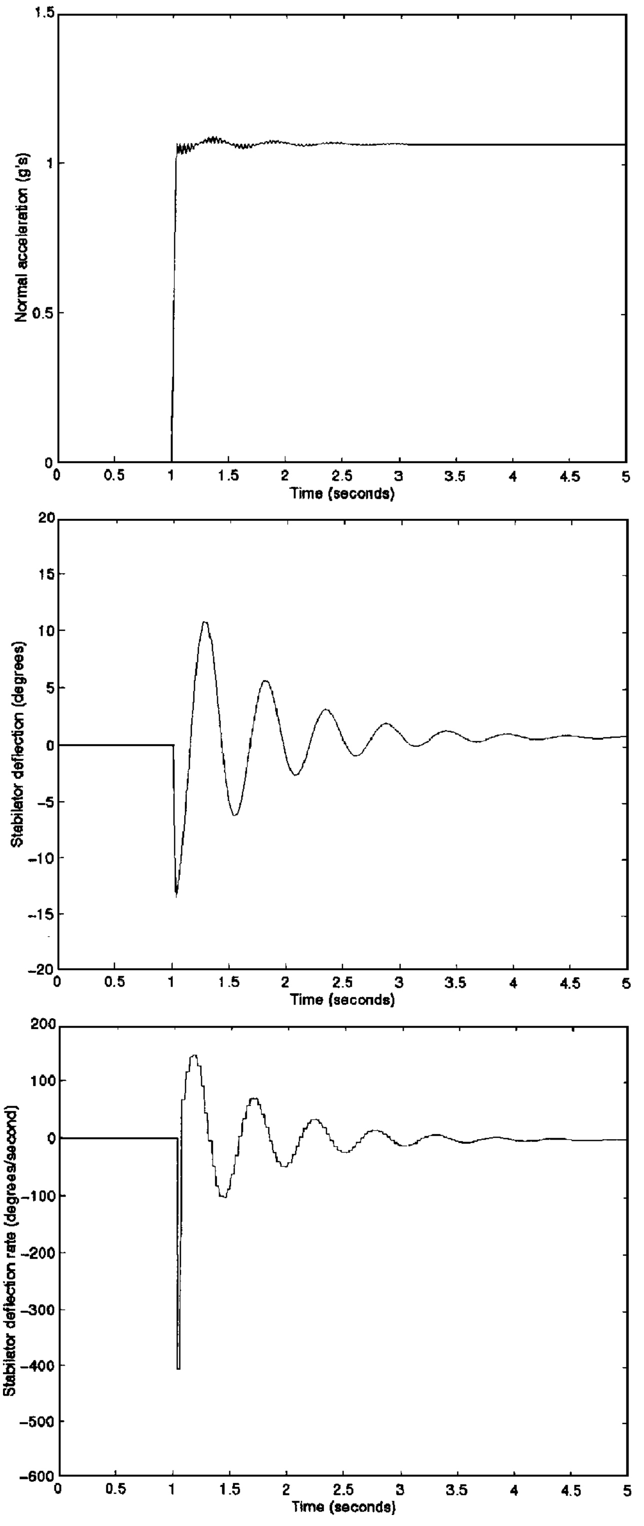

Fig. 4 Unweighted sensitivity step, control, and control rate responses.

control usage and rate of control usage are also shown in Fig. 4. The control usage does not violate the maximum deflection limits, but it is quite large for only a 1-g change in normal acceleration. Because the system is linear, it is easy to see that the maximum deflection limit would be violated for a commanded 2-g change in normal acceleration. The control rate violates the maximum allowable rate limitation, even for a very small command.

The controller found would be undesirable for two reasons: first, the system tracks with a steady-state error; second, the level of performance shown in Fig. 4 is unattainable by the AFTI F-16 due to limitations in the stabilator rate of deflection. Before discussing how to handle these problems, it is important to discuss some objectives the tracking solution should achieve. The following list represents some typical factors that may be important: 1) minimum error to low-frequency commands, 2) no violations of control defection and rate limitations, 3) zero steady-state error to low-frequency commands, 4) minimum overshoot and undershoot, and 5) the quickest response possible given the preceding items.

Item 1 indicates that sensitivity minimization is the proper objective function for ℓ_1 optimization, but items 2–4 indicate that it must be done with certain constraints. Item 5 is inherently built into ℓ_1 optimization for most problems. Methods of incorporating items 2–4 without using sensitivity frequency weights are explored in the sequel. In general, the same problems are encountered in a weighted sensitivity design. The next section tackles item 2. It covers three different approaches for adding control deflection and rate constraints to the error minimization problem.

Control Deflection and Rate Limitations

The preceding section was concerned with solving the one-block problem

$$\inf_{K\,\text{stabilizing}} \|S\|_1 \tag{10}$$

where S is the sensitivity function. This section will first examine the general two-block problem

$$\inf_{K\,\text{stabilizing}} \left\| \begin{matrix} S \\ W_c K S \end{matrix} \right\|_1 \tag{11}$$

where W_c is a weighting on the control usage. The added block in Eq. (11) can be used to ensure that control deflection or rate limitations are not violated.

Because the control rate limitations were violated in the last section, only rate constraints are added in this section. It turns out that once the control rate is properly constrained for the AFTI F-16, the control deflection limitations are not a problem. The ideas presented subsequently, however, easily extend to penalizing control deflections alone or to both control rates and deflections.

To change the second block of Eq. (11) to a penalty on control rate instead of control usage, an appropriate weight must be chosen for W_c. Clearly, the weighting must be chosen so that it effectively takes the derivative of the control signal. This problem is best handled directly in the z domain, with

$$W_c(z) = (z-1)/Tz \tag{12}$$

where T is the sample period. This weighting function, known as the backward Euler transformation, calculates a finite difference gradient between discrete pulses. Because the weighting is in the z domain, the continuous system must be discretized before this weight can be augmented to the problem.

The first approach to solving the rate-constrained tracking problem is to multiply each block in the two-block problem by a desired level of performance. For example, if the one norm of the first block is desired to be less than γ and the maximum control deflection rate is $U_{r_{\max}}$, then the problem becomes

$$\inf_{K\,\text{stabilizing}} \left\| \begin{matrix} (1/\gamma)S \\ \dfrac{1}{U_{r_{\max}}} W_c K S \end{matrix} \right\|_1 \tag{13}$$

If the resulting one norm of this system is less than 1, then both levels of performance are achieved. This follows from the previous assumption that the maximum magnitude of the exogenous input is less than or equal to one. If the goal is to find a solution that has the minimum achievable γ without violating the maximum control rate, this is not the best approach because Eq. (13) would have to be solved iteratively for γ until the resulting system one norm is exactly 1.

A better approach is to solve the following problem:

$$\inf_{K\,\text{stabilizing}} \|S\|_1 \qquad \text{subject to } \|W_c K S\|_1 \le U_{r_{\max}} \tag{14}$$

In Eq. (13), the maximum absolute row sum had to be less than 1 to ensure the one norm of the entire system was also less than 1. In Eq. (14) the individual row sums are separated, with one being minimized while the other is constrained.

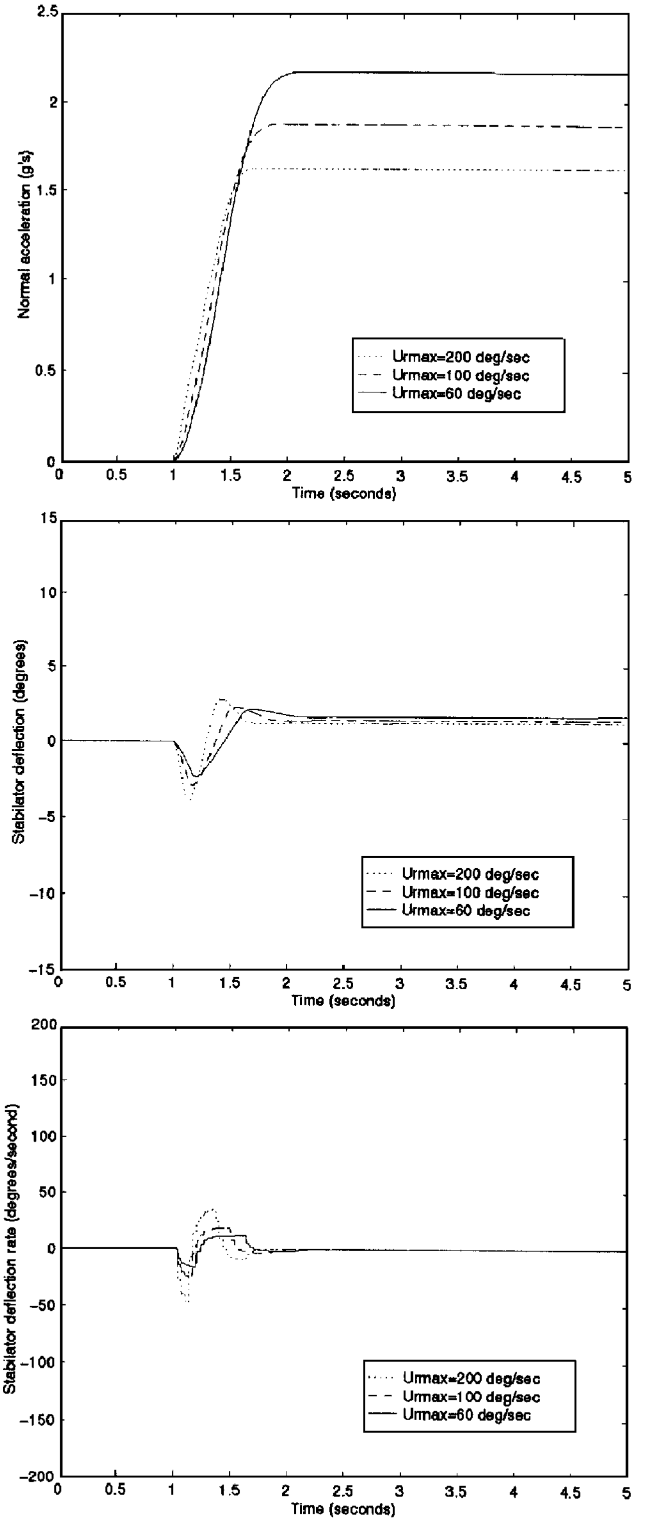

Fig. 5 Unweighted sensitivity step, control, and control rate responses with ℓ_1 constraints on the control rate.

Equation (14) was solved for the AFTI F-16 with control rate limitations of 200, 100, and 60 deg/s. A plot of the system unit step response for all three constraint levels is shown in Fig. 5. The slowest response with the largest steady-state error corresponds to the actual stabilator deflection rate limitation of 60 deg/s. The responses to the other two constraint levels are shown for comparison. This type of plot can also be used for design purposes because it is easy for the designer to see how much performance can be gained if faster control actuators are obtained. Plots of the control deflections and rates are shown in Fig. 5. Notice the control rates are well below their respective ℓ_1 constraints for a unit step input.

The one norm of the objective and the compensator orders are shown in Table 1 for each constraint level. The order of the ℓ_1

Table 1 Comparison of different ℓ_1 constraints on control rate

Constraint, deg/s	One norm	Controller order
200	2.63	35
100	2.88	38
60	3.17	43

Table 2 Comparison of different ℓ_∞ constraints on control rate

Constraint, deg/s	One norm	Controller order
200	2.16	9
100	2.28	13
60	2.43	19

optimal compensators is directly related to the support length of the pulse response, i.e., the number of time steps it takes the pulse response to decay to zero. Because the support length of the pulse response is related to the time it takes the step response to reach steady state, it is easy to see why the controllers found using the described approach have such high orders.

The preceding two approaches imposed ℓ_1 constraints on the control rate. This means that the constraint limitation imposed will not be exceeded for any input into the system bounded in magnitude by 1. Another less conservative option is to ensure that the constraint is not exceeded for a single class of input such as the step command. Unlike the ℓ_1 constraints, these ℓ_∞ types of constraints can only be used on a finite horizon. In other words, they can only be imposed over the support length of the solution. In many cases, however, imposing these constraints over the first few time steps is sufficient.

To understand how a constraint on the step response of the system can be imposed in ℓ_1 optimization, recall the relationship in Eq. (9) between the unit pulse and unit step response. The step response at any particular time step is nothing more than the sum of the pulse response at that time step plus all previous time steps. Therefore, in terms of the pulse response at each time step, these constraints can be imposed with very simple Toeplitz matrices, with ones below the main diagonal and zeros above,

$$\begin{bmatrix} 1 & 0 & \cdots & 0 \\ 1 & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & 0 \\ 1 & 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} \hat{\Phi}(0) \\ \hat{\Phi}(1) \\ \vdots \\ \hat{\Phi}(N) \end{bmatrix} \leq \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} U_{r\max} \tag{15}$$

These constraints can easily be augmented to the constraints in Eq. (4).

The new problem with the augmented step input constraints becomes

$$\inf_{K \text{ stabilizing}} \|S\|_1 \qquad \text{subject to } \|W_c K S w_f\|_\infty \leq U_{r\max} \tag{16}$$

where w_f is a unit step. A plot of the AFTI F-16 normal acceleration step response for control rate constraint levels of 200, 100, and 60 deg/s is shown in Fig. 6. Again, the slowest response with the largest steady-state error corresponds to a constraint of 60 deg/s. Notice that the step responses to this type of constraint are much quicker and have less steady-state error than the ℓ_1 constraints. Control deflections and control rates are also shown in Fig. 6. Now the rates appear to hit the constraint limits, except in the 200-deg/s case. This is due, however, to the finite differencing used on the rates. In each case, the rate constraints are active.

The one norm of the objective and the compensator orders are shown in Table 2 for each constraint level. With this approach, quicker settling times also lead to lower-order controllers. Although all of the step responses in this section meet some constraint level on the control rate, none of them has zero steady-state error. This issue is addressed in the next section.

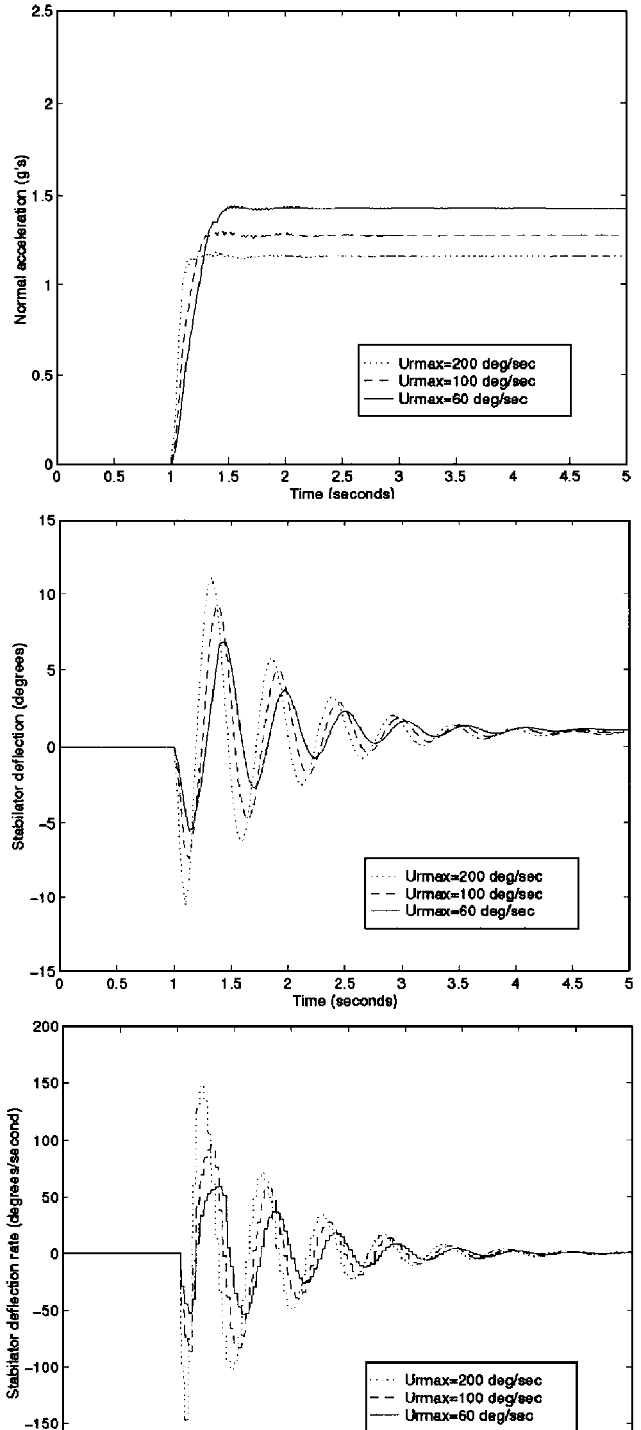

Fig. 6 Unweighted sensitivity step, control, and control rate responses with ℓ_∞ constraints on the control rate.

Steady-State Error and Time-Varying Exponential Weights

Near zero steady-state error to a step input can be enforced by using a weighting on sensitivity. This typically causes problems with control rate usage and overshoot, which then need to be addressed. Thus, no weight will be added to the sensitivity, and zero steady-state error to a step input will be enforced by adding an equality constraint to the ℓ_1 optimization problem. Recall from Eq. (9) that the final value of the step response is simply the summation of the unit pulse response over its entire support length. Therefore, zero steady-state error to a unit step input is guaranteed if the sum of the sampled unit pulse response equals zero. This is not an absolute

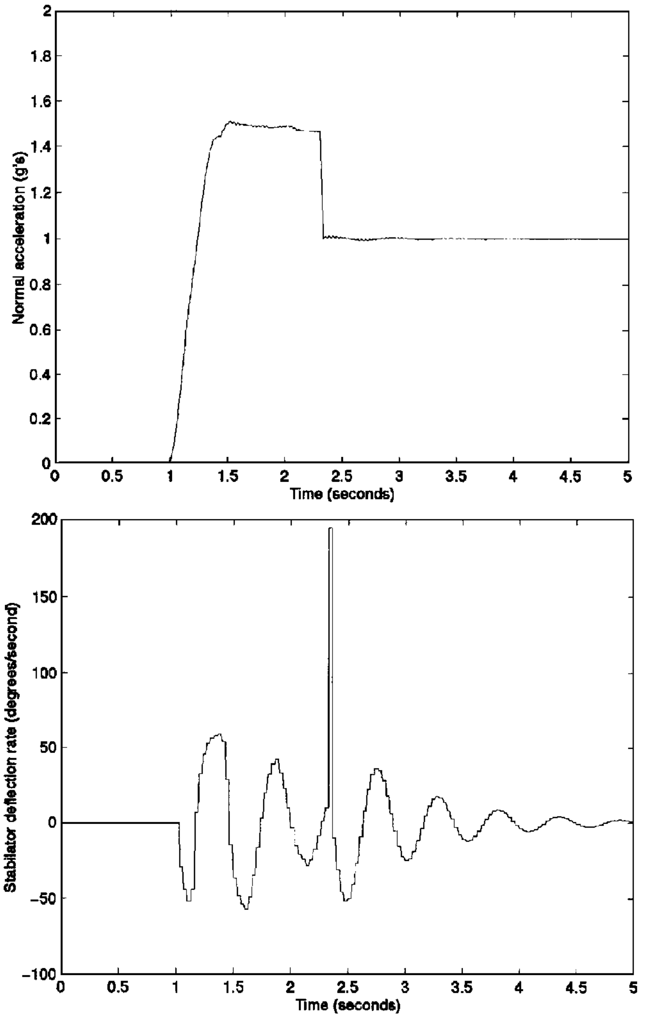

Fig. 7 Unweighted sensitivity step and control rate responses with steady-state error and ℓ_∞ control rate constraints.

summation like the norm calculation; it is simply a summation of the pulse response at each time step. The added equality constraint takes the form

$$[1 \quad 1 \quad \cdots \quad 1]\begin{bmatrix}\hat{\Phi}(0)\\ \hat{\Phi}(1)\\ \vdots \\ \hat{\Phi}(N)\end{bmatrix} = 0 \tag{17}$$

This constraint was added to the problem presented in Eq. (16), with the control rate constraint equal to 60 deg/s. The resulting solution has an objective one norm of 3.00 and the compensator is 44th order. The system response to a 1-g normal acceleration step input is shown in Fig. 7 along with the control rate. Clearly, zero steady-state error is achieved, but the limit on control rate is not met at the nearly discontinuous jump just after 2 s. Notice that the response in Fig. 7 is exactly the same as its counterpart in Fig. 6 up until this jump. This happened because the steady-state error equality constraint was not imposed until the very last time step in the support length. In this system, imposing the constraint any earlier results in a larger one norm, which the optimization rejects. This problem can be overcome with time-varying exponential weights on the norm calculation. Consider multiplying each sampled pulse response by a^{kT}, where k is the sample index, T is the sample period, and $a > 1$. Because this weight gets larger as k gets larger, it effectively penalizes late errors over early ones.

The preceding system was rerun with a time-varying exponential weight added to the norm calculation. With $a = 1.1$, the objective one norm was 3.59 and the controller was 26th order. A plot of the system response to a step input is shown in Fig. 8. Adding the exponential weights increased the one norm as expected. The weighting also decreased the settling time and thus the controller order. The step response at this point now meets all of the criteria established, except the overshoot issue. This problem is discussed in the next section.

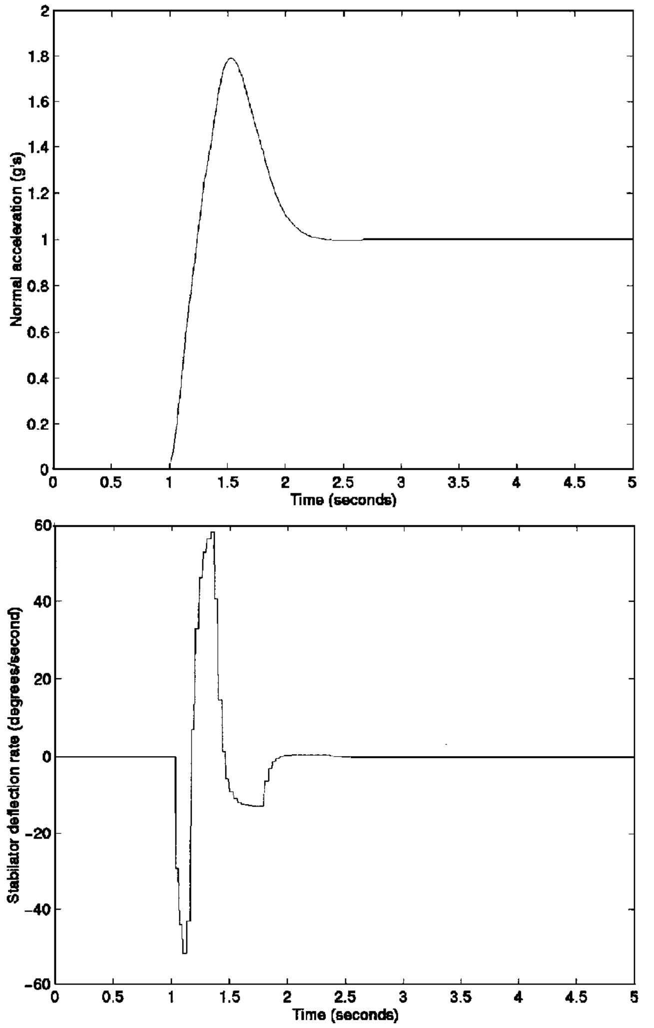

Fig. 8 Unweighted sensitivity step and control rate responses with steady-state error, ℓ_∞ control rate constraints, and time-varying exponential weights.

Overshoot and Undershoot Limitations

Problems with excessive overshoot and undershoot in the step response can be handled in exactly the same manner as excessive control deflections and rate violations. To demonstrate this concept, a very small reduction is done on the overshoot for the step response shown in Fig. 8. Theoretically, the overshoot can be reduced to any desired level at the expense of a slower response. However, it was extremely difficult to find a solution for the problem presented with the current ℓ_1 optimization software. The multiblock problem contains so many constraints and delays that calculation attempts alone are extremely expensive in terms of computer time. Further, the linear programming routine in the software has difficulty solving very large systems of equations, possibly due to scaling problems. The system response in Fig. 8 has an overshoot of about 80%. An ℓ_∞ type constraint on the overshoot was added to the problem to ensure that the overshoot would be less than 70% to a step input. The resulting solution had an objective one norm of 3.40, and the controller was 28th order. A plot of the step response with the added constraint is shown in Fig. 9. Figure 10 shows a Bode plot of the controller. Note the high gain at low frequency to achieve tracking, as well as the lag-lead behavior needed to keep the control usage down and respond quickly.

This step response is still less than ideal; however, all of the tools to shape and constrain the response are now available. As the ℓ_1 optimization software becomes more reliable and efficient, a designer

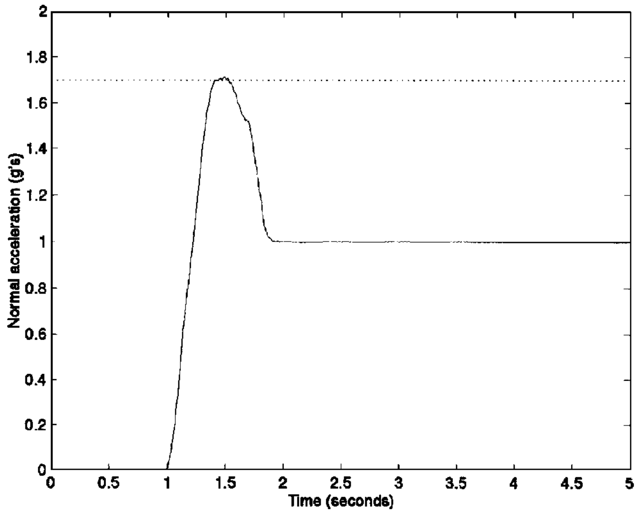

Fig. 9 Unweighted sensitivity step response with control rate, steady-state error and overshoot constraints, and time-varying exponential weights.

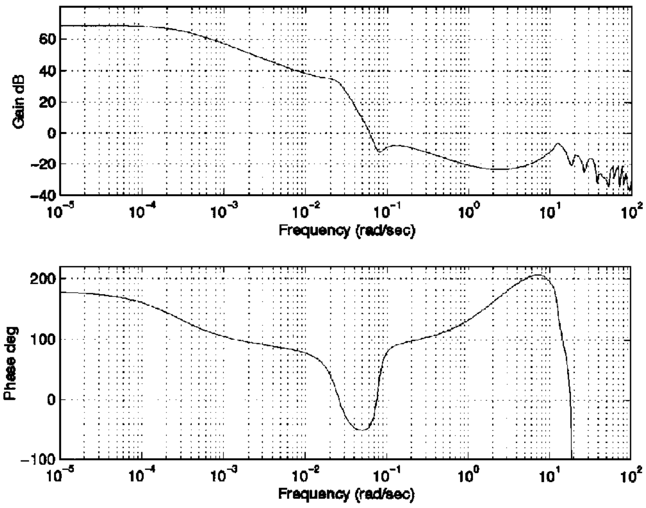

Fig. 10 Bode plot for the controller of Fig. 9.

should be able to use all of the techniques presented up to this point to find a compensator that meets all of the tracking requirements. Unfortunately, this compensator may be extremely high order and therefore impractical to use. Note that the designs shown up to this point are not intended to represent the best ones possible. Rather, they are shown to clearly illustrate the capabilities of ℓ_1 optimization. The next section on model matching demonstrates one way to use ℓ_1 optimization and still produce controllers at or about the order of the original discrete system. Some of the problems encountered in this final design will be handled with alternatives to the methods shown up to this point.

Model Matching

One way to counter order inflation in ℓ_1 optimization is to solve a one-block problem instead. Because these problems can be solved exactly, without delay augmentation, the resulting controllers tend to be much smaller (usually about the order of the unweighted plant). An added benefit to using one-block systems is that they can be solved much faster and more reliably than multiblock systems with the current ℓ_1 optimization software.

In many cases, however, the constraints discussed in the preceding sections cannot be imposed with one-block systems. Therefore, a one-block system must be chosen that incorporates as many of the design criteria required for good tracking as possible. One way to accomplish this objective is to model match the design problem to a system that has the desired tracking characteristics.

A model-matching design for the SISO AFTI F-16 problem that has been discussed throughout is shown in Fig. 11. A small penalty on control usage, similar to the one discussed earlier, was added to the system to make the resulting ℓ_1 problem nonsingular. In Fig. 11, H is the ideal closed-loop model given in continuous time by $H(s) = 4/(s+4)$. This closed-loop model was chosen because its step response is relatively quick and has no overshoot and no steady-state error.

The one norm of the solution to this design problem is 0.37 and the compensator is fifth order. A plot of the AFTI F-16 step response to a commanded 1-g normal acceleration change is shown in Fig. 12, along with the response of the ideal model. The step response of the solution is approximately 0.37 g larger than the step response of the ideal model at steady state. This is because ℓ_1 optimization penalizes the maximum error and does not address steady-state error in this setup. To make the system response match the ideal model's response, the commanded normal acceleration has to be multiplied by a gain. This gain is the reciprocal of the dc gain of the discrete closed-loop transfer function,

$$T_{yr}(z) = \frac{K(z)G(z)}{1 + K(z)G(z)} \tag{18}$$

For this problem, the gain equals 0.73.

Notice that with this particular approach there is no direct way to ensure the preceding solution will not violate control deflection and rate limitations. In this problem, the control rate limits were violated in the first few time steps. However, the closed-loop system still performs well if the step input is first passed through a continuous prefilter equal to

$$F(s) = 10/(s + 10) \tag{19}$$

and a rate limiter is added to the control signal. With the added prefilter, the system sees a smooth continuous approximation of a step input, rather than a discontinuous step input. The new input is actually a more realistic representation of a pilot command.

The system was tested with the prefilter, gain adjustments on the input, and a control rate limiter set at ±60 deg/s. The response is shown in Fig. 13. This response has no overshoot, no steady-state error and was achieved with a fifth-order controller and a small gain on the input. A Bode plot of the resulting controller is shown in Fig. 14. This controller does not insert high gain at low frequency,

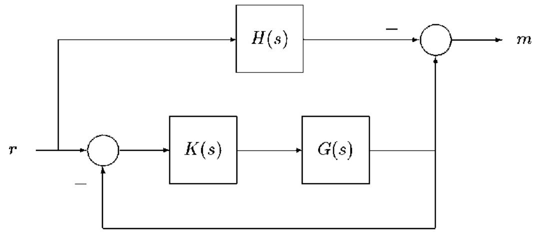

Fig. 11 Closed-loop model matching diagram.

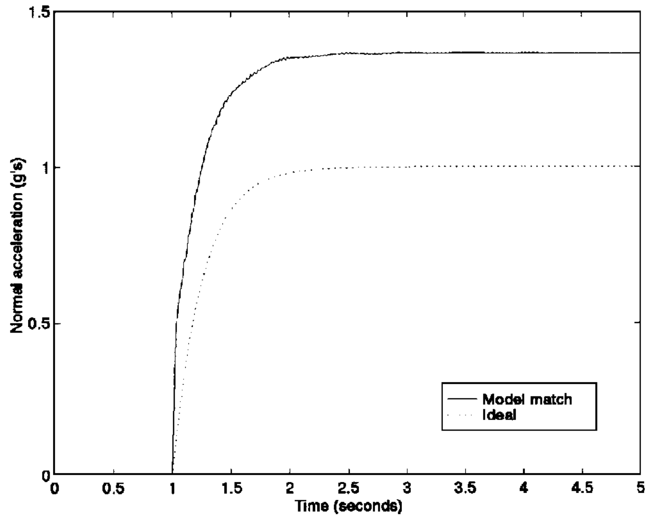

Fig. 12 Model-matching step response.

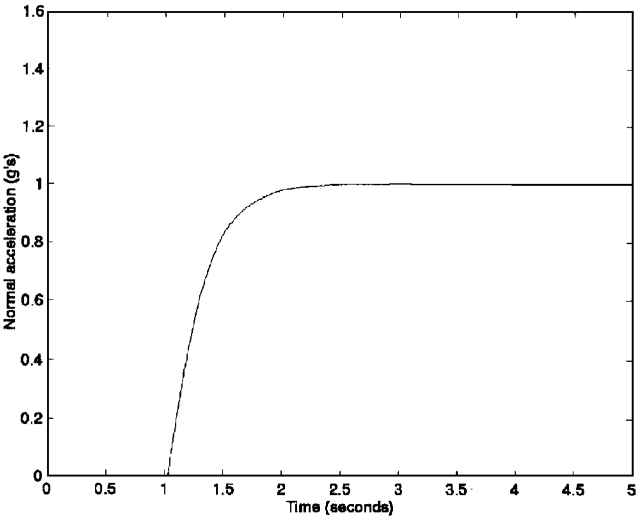

Fig. 13 Model-matching step response with scaling factor, prefilter, and control rate limiter.

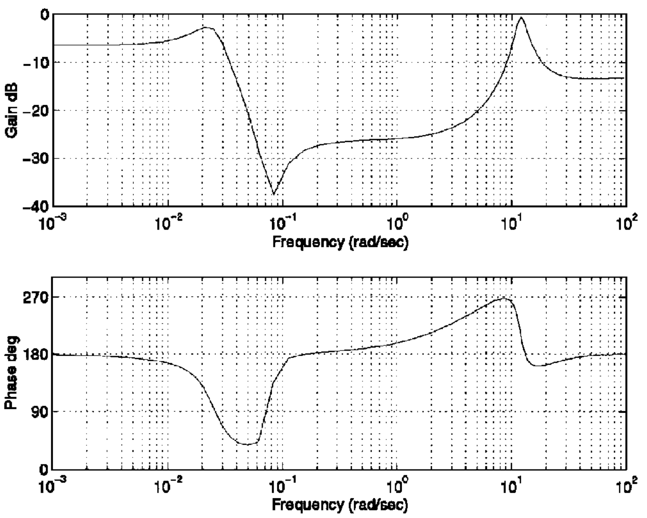

Fig. 14 Bode plot for the controller of Fig. 13.

because zero steady-state error is not enforced by the controller, but rather by the command gain. This controller also has a lag-lead structure, as in the constrained sensitivity minimization case. The vector gain margins (VGM) in decibels and vector phase margins (VPM)[7] in degrees of this system are

$$-11.4 \le \text{VGM} \le 10.4, \qquad \text{VPM} = \pm 42.9$$

These margins are very high because of a good match to an ideal model with a good loop shape.

Conclusions

Several methods of using ℓ_1 optimization to solve tracking problems were presented. Specifically, constraints that could be added to the standard ℓ_1 problem to handle control deflection and rate limitations, zero steady-state error requirements, and overshoot limitations were discussed. Although these constraints theoretically allow the control designer to tailor a system's time response, they tend to lead to multiblock problems and, thus, high-order controllers. These highly constrained multiblock problems also tend to be computationally expensive and difficult to solve with the current solution techniques. If the control designer can restrict the design to a one-block problem, however, controller orders at or near the order of the weighted plant are possible. This is because one-block ℓ_1 optimization problems can be solved exactly in this case without DA. The utility of doing a one-block ℓ_1 optimization design was demonstrated with a simple model-matching problem.

Appendix: Aircraft Data

The aircraft design model used consists of an actuator servo G_a and the linearized longitudinal equations of motion for the aircraft G_p, referred to as the core plant. The state-space representation of the continuous system is found by concatenating the two components.

The four states in the longitudinal model are forward speed (u in feet per second), angle of attack (α in radians), pitch angle (θ in radians), and pitch rate (q in radians per second). The input to G_p is the stabilator deflection (δ_e in radians) and the output is the normal acceleration (n_z in g). G_p is given by

$$\begin{bmatrix} \dot{u} \\ \dot{\alpha} \\ \dot{\theta} \\ \dot{q} \end{bmatrix} =$$

$$\begin{bmatrix} -1.485e-2 & 3.738e+1 & -3.220e+1 & -1.794e+1 \\ -8.000e-5 & -1.491e+0 & -1.300e-3 & 9.960e-1 \\ 0.000e+0 & 0.000e+0 & 0.000e+0 & 1.000e+0 \\ -3.600e-4 & 9.753e+0 & 2.900e-4 & -1.904e+1 \end{bmatrix}$$

$$\times \begin{bmatrix} u \\ \alpha \\ \theta \\ q \end{bmatrix} + \begin{bmatrix} 2.140e-3 \\ -1.880e-1 \\ 0.000e+0 \\ -1.904e+1 \end{bmatrix} \delta_e$$

$$n_z = [1.500e-3 \quad 3.5264e+1 \quad 2.720e-2 \quad -3.340e-1]$$

$$\times \begin{bmatrix} u \\ \alpha \\ \theta \\ q \end{bmatrix} + [-4.366e+0]\delta_e$$

The input to G_a is the commanded stabilator deflection (δ_{ec} in radians), and the output is the stabilator deflection. G_a is given by

$$\dot{x}_a = [-2.000e+1]x_a + [2.000e+1]\delta_{ec}$$

$$\delta_e = [1.000e+0]x_a + [0.000e+0]\delta_{ec}$$

The discrete aircraft plant G equals $G_a G_p$ discretized at 30 Hz using a ZOH.

Acknowledgments

This work was sponsored in part by the U.S. Air Force Office of Scientific Research. The authors would like to thank Ignacio Diaz-Bobillo for allowing them to use a copy of his MATLAB ℓ_1 software for this work.

References

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

[2] Dahleh, M. A., and Diaz-Bobillo, I. J., *Control of Uncertain Systems*, Prentice–Hall, Englewood Cliffs, NJ, 1995.

[3] Vidyasagar, M., "Optimal Rejection of Persistent Bounded Disturbances," *IEEE Transactions on Automatic Control*, Vol. AC-31, No. 6, 1986, pp. 527–534.

[4] Dahleh, M. A., and Pearson, J. B., "ℓ_1 Optimal Feedback Controllers for MIMO Discrete-Time Systems," *IEEE Transactions on Automatic Control*, Vol. 32, No. 4, 1987, pp. 314–327.

[5] Youla, D. C., Jabr, H. A., and Bongiorno, J. J., "Modern Wiener-Hopf Design of Optimal Controllers, Part II: The Multivariable Case," *IEEE Transactions on Automatic Control*, Vol. AC-21, 1976, pp. 319–338.

[6] Maciejowski, J. M., *Multivariable Feedback Design*, Addison–Wesley, Reading, MA, 1989, Chap. 6.

[7] Franklin, G. F., Powell, J. D., and Emami-Naeini, A., *Feedback Control of Dynamic Systems*, 2nd ed., Addison–Wesley, Reading, MA, 1991, Chap. 5.