JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 23, No. 3, May–June 2000

Loop Shaping Design for Missile Autopilots: Controller Configurations and Weighting Filter Selection

Pablo A. Iglesias*
Johns Hopkins University, Baltimore, Maryland 21218
and
Thomas J. Urban†
Johns Hopkins University, Applied Physics Laboratory, Laurel, Maryland 20723-6099

The design of a homing, tail-controlled missile autopilot using a loop shaping approach is considered. The designer input in this control methodology consists of a set of shaping filters that in turn lead to an $\mathcal{H}_\infty$ optimal controller that guarantees stability. Several configurations are contrasted. These differ as to the location in the loop of the shaping filters and optimal controller. It is shown that the choice of configuration has implications as to the form of the shaping filters. It is also demonstrated how intuition from classical three-loop controllers can be used to select these filters. Finally, for one configuration, analytic formulas for a set of generic shaping filters are derived that lead to good tracking performance and robustness margins.

I. Introduction

THE classical approach to autopilot design for homing interceptor missiles has remained basically unchanged since the Sparrow designs of the early 1950s.[1,2] Three single-input/single-output controllers are designed to control roll angle, as well as pitch and yaw angles. The goal of the roll loop is to prevent the rolling of the missile to minimize coupling of guidance signals between yaw and pitch. The pitch and yaw angle controllers keep the airframe dynamically stable throughout its flight envelope. This envelope includes a large range of operating conditions, which vary according to altitude, angles of attack, and Mach number.

The three-loop approach to autopilot design works adequately whenever the interaction between channels (primarily the yaw and pitch loops) is minimal. As the side-slip and attack angles increase, however, the coupling between the channels increases, resulting in poor performance. For this reason, a multivariable approach to autopilot design is desirable.

In addition to minimizing the dynamic coupling between the different loops, a further requirement of a controller is that it covers as much of the flight envelope as possible. The standard approach to designing controllers for the different operating regions is to schedule the gains of the three-loop topology controller. This is a fairly straightforward procedure. The suitability of the controller for gain scheduling is an important design requirement. For this reason, robust controllers such as those obtained via $\mathcal{H}_\infty$ optimization techniques prove useful. Because a robust controller ensures satisfactory performance around an operating point, fewer linear designs are needed.

In this paper we consider the design of a missile autopilot using the $\mathcal{H}_\infty$ loop shaping procedure (LSP) introduced by McFarlane and Glover.[3] Since its introduction, the LSP has received considerable attention as a means of designing robust multivariable controllers.[4–10] In particular, Hyde and Glover conducted a comprehensive design study for the longitudinal controller of a vertical/short takeoff and landing (VSTOL) aircraft model leading to successful flight tests on a research Harrier.[11]

Received 4 June 1999; presented as Paper 99-3973 at the AIAA Guidance, Navigation, and Control Conference, Portland, OR, 9–11 August 1999; revision received 8 November 1999; accepted for publication 11 November 1999.

*Associate Professor, Department of Electrical and Computer Engineering, 3400 North Charles Street; pi@jhu.edu.

†Supervisor, Preliminary Design and Algorithm Development Section, Dynamic Systems Group, Air Defense Systems Department, 11100 Johns Hopkins Road; thomas.urban@jhuapl.edu.

Several of the design requirements of the VSTOL aircraft are similar to that of the missile autopilot. The VSTOL aircraft is highly coupled, a coupling arising from pitch moments generated by changes in the thrust nozzle setting. Moreover, the operation of a VSTOL aircraft ranges over differing regimes, ranging from the hover mode, where the plane is unstable, to forward flight. As in the missile autopilot, this imposes a gain scheduling requirement. These similarities suggest that it may be possible to use the LSP to design practical controllers for the missile autopilot.

The designer input in the LSP comes through the choice of shaping filters on the plant. In this paper we demonstrate how different possible controller configurations lead to different choices in the form of the weighting filters. We contrast different controller configurations and provide analytic formulas for the shaping filter in one particular configuration. Related work on the use of the LSP for missile autopilots is found in the literature.[8,12,13]

The rest of this paper is organized as follows. We begin by providing some details of the LSP. We then introduce the classical autopilot topology and outline the requirements on the controller. The different possible autopilot topologies are contrasted, and controllers are designed for each of the possible topologies. Finally, we offer analytic expressions for a particular shaping filter arrangement that offers good performance and robustness.

One aspect of the design process that will not be considered here is the scheduling of the controllers. The gain scheduling of controllers obtained using the LSP has been considered in detail for the VSTOL aircraft.[11] Note that this discussion is not specific to the particular aircraft but to the design procedure and, hence, applies to the missile autopilot.

II. LSP

Although $\mathcal{H}_\infty$ control theory provides a powerful means of designing robust linear controllers, it is the particular transfer function whose $\mathcal{H}_\infty$ norm is minimized that has the greatest impact on the suitability of the resulting controller.

Early choices for the optimization problem included weighted-sensitivity[14] and mixed-sensitivity[15] minimization. In both of these cases, weights are used to penalize closed-loop performance and robustness objectives. Whereas these weights are used to influence the closed-loop transfer function by means of the optimization problem, the weights themselves do not appear in the final closed-loop system.

McFarlane and Glover introduced an alternative approach to $\mathcal{H}_\infty$ controller design.[3] This scheme is based on a robust stabilization problem (RSP) in which the closed-loop system's tolerance to uncertainty is maximized. This robustness margin is measured in terms of the size of the allowable perturbations to a normalized left-coprime factorization of the nominal plant. In particular, suppose that a linear