Robust H-Infinity Control of a Steerable Marine Radar Tracker

Daniel J. Auger, Stuart Crawshaw
Stephen L. Hall

BAE Systems Integrated System Technologies Ltd,
Broad Oak Works, Airport Service Road, Portsmouth, Hants, PO3 5PQ,
United Kingdom (Tel: +44 2392 515350; e-mail: stuart.crawshaw@baesystems.com)

Abstract: This paper describes the application of a robust control technique to a steerable marine radar tracker intended to provide good performance with minimum operator intervention over the course of its lifecycle. The paper shows that the sightline steering problem can be decoupled from the target observation problem and uses a well-known robust control technique (H_∞ loop-shaping) to synthesize a controller. Analytic bounds on the stability of the closed loop are stated by considering a model set parameterized on Vinnicombe's ν-gap metric. Experimental verification exercises are briefly described, and proposals for formal validation work using interpolation techniques are made.

1. INTRODUCTION

This paper describes an application of robust control to a radar tracker series used on a marine platform deployed in varied environmental conditions during a several-year operating life. Earlier models used an arrangement of analogue feedback loops and carefully-tuned 'feedforward' terms. A careful set-up procedure was required for each tracker, and daily adjustment was required to maintain the equipment on specification. A solution requiring little operational maintenance is preferred. The tracker must follow a variety of airborne targets, some of which may be manoeuvring. A customer requirement specifies that the radar must be capable of tracking targets with stated steady-state angular velocities and accelerations; for commercial reasons, the authors are unable to give values. The acceleration requirement is new and earlier solutions only met the velocity requirement. The tracker operates in safety-critical situations.

The tracker geometry is shown in Fig. 1. A sensor suite is mounted such that it can be steered in training and elevation; data from the sensors is merged in an algorithm based on an extended Kalman filter. The steering mechanism consists of a pair of pulse-width modulated electric motors and a mechanical gearbox and internal electronics are arranged such that a constant input voltage signal will result in a constant steady-state angular velocity.

For the new design, the authors have chosen to use H_∞ loop-shaping (McFarlane and Glover, 1992) a robust control technique combining H_∞ optimization with traditional frequency-response techniques. This method is documented in several widely available control textbooks (Glad and Ljung, 2000; Vinnicombe, 2000), and has been successfully applied to a flight control problem by a U.K. research agency (Hyde *et al*, 1995). The authors chose this technique because a robust method should reduce the need for frequent re-tuning and allow one control algorithm to be used for all production model trackers.

Sec. 2 of this paper introduces relevant ideas from the literature; sec. 3 describes the application to the authors' problem.

1.1 Notation

The notation used in this paper follows the pattern of many in standard control texts, e.g. Zhou *et al* (1996). L_∞ is the Banach space of matrix-valued functions that are essentially bounded on the unit circle; the corresponding norm is

$$\left\| Q \right\|_\infty := \operatorname*{ess\,sup}_{\Omega \in (-\pi, \pi]} \bar{\sigma}[Q(e^{j\Omega})]$$

where $\bar{\sigma}(X)$ denotes the maximum singular value of the matrix X. H_∞ is a subspace of L_∞ containing those functions which are analytic and bounded outside the unit disc; the corresponding norm is

$$\left\| Q \right\|_\infty := \sup_{|\Omega| > 1} \bar{\sigma}[Q(e^{j\Omega})] = \sup_{\Omega \in (-\pi, \pi]} \bar{\sigma}[Q(e^{j\Omega})].$$

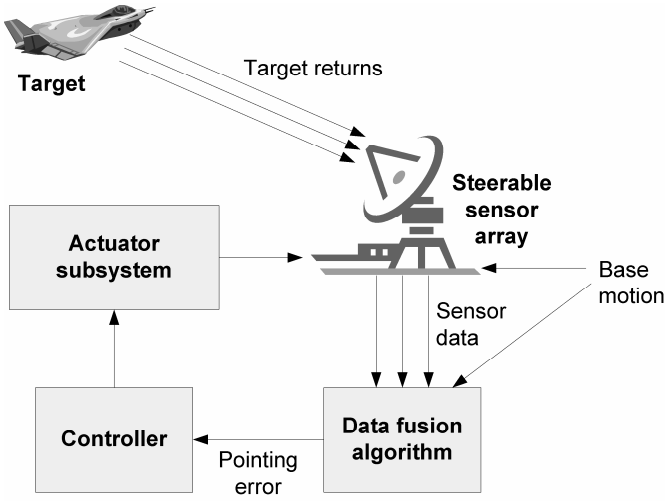

Fig. 1. Tracker schematic. This diagram gives a high-level conceptual illustration of the tracker configuration.

Effectively, H_∞ is the 'stable' part of L_∞. RH_∞ is the subspace of H_∞ comprising all functions in H_∞ that can be expressed as proper rational transfer functions with real-valued coefficients. $b(P,C)$ is the normalized coprime factor stability margin obtained when plant P is stabilized with controller C (see sec. 2.2). $\delta_\nu(P,P_1)$ is the ν-gap between systems P and P_1 (see sec. 2.3). z represents the inverse of the unit delay operator. Given a function $G(s)$, wno(G) denotes the winding number evaluated on a contour indented into the complex left half plane around any imaginary poles of G. $\eta(G)$ is the number of open right half plane poles of $G(s)$. Dimensions will be omitted wherever the sense is obvious from the context.

2. Theoretical Background

2.1 Observers and Separation Theory

Observers (a.k.a. estimators) are described in many standard texts e.g. Franklin *et al* (1994; 1998). Observers are models of a system's dynamics, driven by a combination of known inputs and output feedback, that is used to estimate the system state to be used in state feedback. A typical configuration for a state-feedback control system using discrete-time observer is shown in Fig. 2. The matrix K is a constant gain, not a dynamic compensator. An important property arising from this configuration is the *separation principle*: the system of Fig. 2. is governed by the matrix equation

$$\begin{bmatrix} x(k+1) \\ \hat{x}(k+1) \\ \hline y(k) \end{bmatrix} = \left[\begin{array}{cc|c} F & -GK & G_r \\ LH & F-GK-LH & 0 \\ \hline H & 0 & 0 \end{array}\right] \begin{bmatrix} x(k) \\ \hat{x}(k) \\ \hline r(k) \end{bmatrix}$$

It is well-known that substituting the observation error $e(k) \coloneqq x(k) - \hat{x}(k)$ yields

$$\begin{bmatrix} x(k+1) \\ e(k+1) \\ \hline y(k) \end{bmatrix} = \left[\begin{array}{cc|c} F-GK & GK & G_r \\ 0 & F-LH & G_r \\ \hline H & 0 & 0 \end{array}\right] \begin{bmatrix} x(k) \\ e(k) \\ \hline r(k) \end{bmatrix}$$

Because of the block-triangular structure of the state transition matrix partition, the poles of the closed-loop system are the solutions to

$$\det(\lambda I - [F - GK]) \cdot \det(\lambda I - [F - LH]) = 0$$

Thus, the poles of the system are the poles that would appear with true state feedback together with the observer poles: the state observation stability problem is decoupled from the control design problem. Both elements will have an effect on the system's initial response to a disturbance r, but the stability problems are effectively separate

A Kalman filter is a particular form of observer using the structure outlined above to generate optimal (in an l_2-norm sense) estimates of the measured quantities (Welch and Bishop, 1995). As far as this paper is concerned, the Kalman filter is seen simply as an observer with a varying gain parameter L.

2.2 Robust Control and H_∞ Loop-Shaping

Robust control uses formal mathematics to design controllers that remain functional in the presence of a specified level of uncertainty. The formulation used in this paper is based on a *normalized right-coprime factorization* (Fig. 3). Given P, factors N and M are chosen such that they satisfy $P = NM^{-1}$ and $\tilde{M}M + \tilde{N}N = I$: the perturbed system $P_1 = (N + \Delta_N)(M + \Delta_M)^{-1}$. Such structures represent uncertainty in the form of "low-frequency parameter errors, neglected high-frequency dynamics and uncertain [outside-unit disk] poles & zeros"

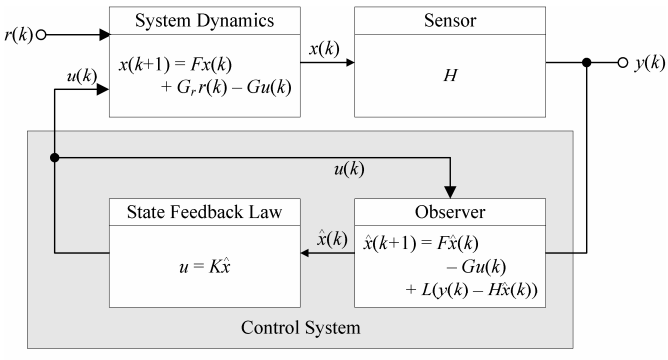

Fig 2. State-feedback control system with an observer. The observer attempts to estimate the state of the system dynamics; the estimated state is used to approximate a state feedback control law.

(Zhou *et al*, 1996). It is possible to define the *stability margin* when stabilized with a compensator C as

$$b(P,C) := \begin{cases} \left\| \begin{bmatrix} P \\ I \end{bmatrix} (I + CP)^{-1} \begin{bmatrix} C & I \end{bmatrix} \right\|_\infty^{-1} & \text{if the loop is stable,} \\ 0, & \text{otherwise.} \end{cases}$$

This definition may be found in many standard works on robust control (Vinnicombe, 2000; Zhou et al, 1996). The same works also contain H_∞ synthesis techniques for designing a controller to maximise $b(P,C)$; commercial software packages are available to perform much of this work (Balas *et al*, 2007).

A drawback of many robust control methods is that it is possible to achieve robust stabilization with good nominal performance or robust performance with good nominal stabilization but not both at the same time. H_∞ loop-shaping (McFarlane and Glover, 1992; Vinnicombe, 2000; Zhou *et al*, 1996) attempts to circumvent this by allowing the control designer to use traditional frequency-response compensation techniques for to specify performance, then 'robustify' the system using H_∞ optimization. The procedure has the following stages:

1. choose 'weights' W_i, W_o such that the weighted feedback-path frequency response $P_W = W_\text{o} P W_\text{i}$ will give the required steady-state error and noise rejection properties (variables of equal importance should be weighted so that they are of comparable size);
2. use H_∞ optimization to generate a controller C_∞ maximising (or, more usually, nearly maximising) $b(P_W, C_\infty)$ – in one source, this step is described as 'robustification';
3. form C, the controller for P, from $C = W_\text{i} C_\infty W_\text{o}$.

In order to achieve good robust performance it is necessary to ensure that $b(P_W, C_\infty)$ is large relative to the expected level of uncertainty. Rules of thumb are often quoted but hard to find in the literature: achieving a $b(P_W, C_\infty)$ value in the range 0.3–0.4 is often considered to indicate good robustness.

2.3 Vinnicombe's ν-Gap Metric

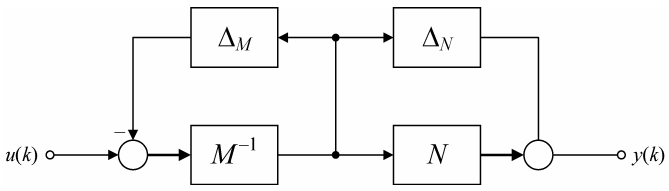

Fig. 3. Normalized right-coprime factorization. This structure can be used to represent a system with uncertain dynamics: the delta blocks are perturbations to the nominal system.

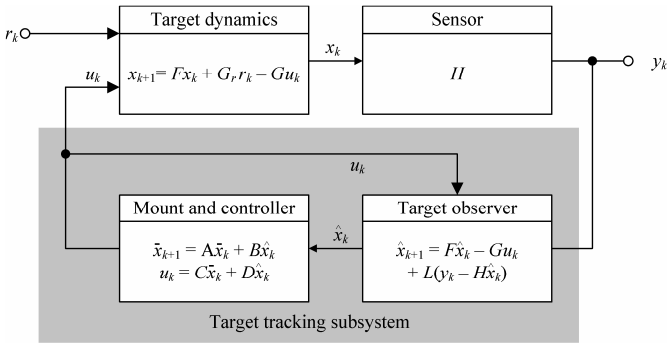

Fig. 5. Separation of tracking and steering (sec. 3.2).

The ν-gap metric (Vinnicombe, 2000) measures the difference between two systems in a way directly comparable with the normalized coprime factor stability margin $b(P, C)$. The notation $\delta_\nu(P, P_1)$ is used to denote the ν-gap between system P and P_1, defined as

$$\delta_\nu(P, P_1) \equiv \delta_\nu(P_1, P) := \inf_{\Delta \in \mathbf{F}} \|\Delta\|_\infty$$

where

$$\mathbf{F} = \left\{ \Delta = \begin{bmatrix} \Delta_N \\ \Delta_M \end{bmatrix} : \begin{array}{c} (N + \Delta_N)(M + \Delta_M)^{-1} = P_1, \\ \|\Delta\|_\infty \in \mathcal{L}_\infty, \\ \text{wno det}\,(M + \Delta_M) = \eta(P_1) \end{array} \right\}.$$

Effectively, the ν-gap between two systems is the smallest perturbation in L_∞ to a normalized right coprime factorization that would make one system equal to another while preserving the number of open-loop unstable poles. The detailed mathematics is beyond the scope of this paper, but an important result can be obtained: given two systems P, P_1 and a compensator C,

$$\arcsin b(P_1, C) \geq \arcsin b(P, C) - \arcsin \delta_\nu(P, P_1)$$

The ν-gap indicates the worst-case degradation in stability margin that one would experience when using a perturbed plant with a controller designed for a nominal one. When using H_∞ loop-shaping, it is the *weighted* ν-gap δ_ν (P_W, $W_\text{o} P_1 W_\text{i}$) that is relevant, since the performance is stated in terms of P_W.

3. Application to the Radar Tracker

3.1 System Configuration

The system from a control perspective (Fig. 4) comprises five elements: the target dynamics (relative to the tracker); the sensor array; the Kalman observer; the tracker dynamics; and a control law. Apart from the control law, all aspects of the system are 'givens': the observer has been inherited, and though it would make sense to take a holistic approach to the

tracking problem if starting afresh it was uneconomic to re-visit its design in the present programme. The sensor dynamics are accounted for in the Kalman filter so for the sake of simplicity they will be ignored in this paper. The tracker kinematics – the angular positions and velocities of the sensor array – are measurable: these are fed-back into the observer in exactly as they affect the real-world tracking dynamics.

3.2 Separation of Target Tracking and Steering

A simple block diagram showing the feedback loop in terms of state-space dynamics is shown in Fig. 5; for convenience, the dynamics of the steering dynamics and the controller have been combined. Using the notation of the diagram, the system is governed by the dynamics

$$\begin{bmatrix} x(k+1) \\ \bar{x}(k+1) \\ \hat{x}(k+1) \\ \hline y(k) \end{bmatrix} = \left[\begin{array}{ccc|cc} F & 0 & 0 & G_r & -G \\ 0 & A & B & 0 & 0 \\ LH & 0 & F-LH & 0 & -G \\ \hline H & 0 & 0 & 0 & 0 \end{array}\right] \begin{bmatrix} x(k) \\ \bar{x}(k) \\ \hat{x}(k) \\ \hline r(k) \\ u(k) \end{bmatrix}$$

together with the feedback law

$$u(k) = C\bar{x}(k) + D\hat{x}(k).$$

It is easy to eliminate $u(k)$, and by defining a change of state variables $e(k) = x(k) - \hat{x}(k)$ the following is obtained:

$$\begin{bmatrix} x(k+1) \\ \bar{x}(k+1) \\ e(k+1) \\ \hline y(k) \end{bmatrix} = \left[\begin{array}{ccc|c} F-GD & -GC & GD & G_r \\ B & A & -B & 0 \\ 0 & 0 & F-LH & G_r \\ \hline H & 0 & 0 & 0 \end{array}\right] \begin{bmatrix} x(k) \\ \bar{x}(k) \\ e(k) \\ \hline r(k) \end{bmatrix}.$$

Since the state transition matrix is block diagonal, the poles of this system are the solution to the eigenvalue equation

$$\det\left(\lambda I - \begin{bmatrix} F-GD & -GC \\ B & A \end{bmatrix}\right) \cdot \det(\lambda I - [F-LH]) = 0$$

The poles are those that one would get if the controller was driven by the true target state, together with those one would get from the observer if it were not connected in feedback. This is a powerful result, since it states that the stability of the steering process is independent of the stability of the observation process: this is not intuitive, and holds because the output of the process is available for use in the observer.

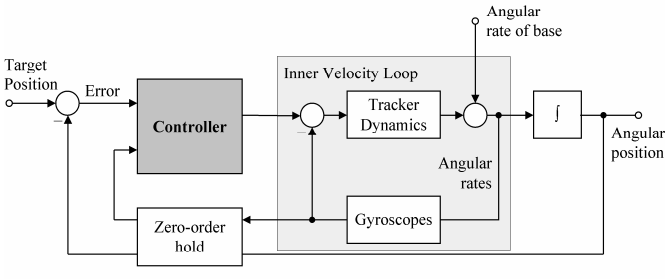

Fig. 6. Block diagram for control design (sec. 3.3).

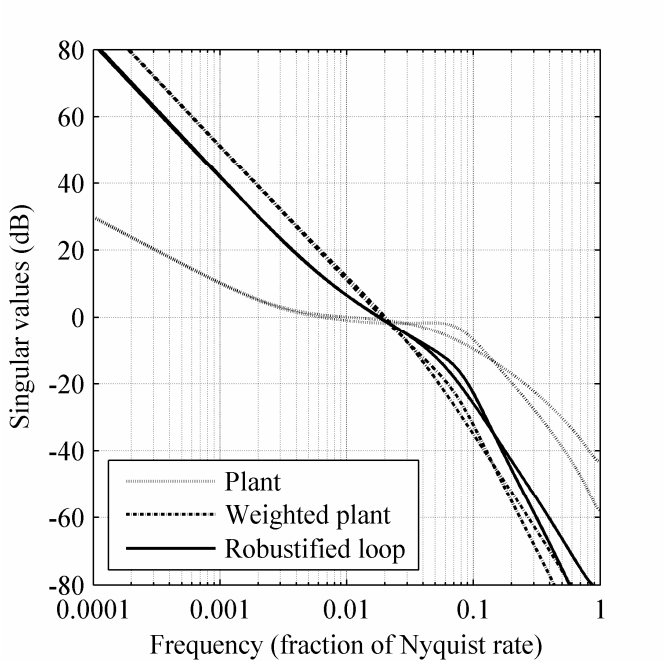

Fig. 7. Singular values (frequency response) of plant before and after weighting, and after loop robustification using H_∞ synthesis.

One way of interpreting this is to note that the problem is effectively one of reduced-order observer design. A system with a full-order observer would have the structure:

$$\begin{bmatrix} x(k+1) \\ \bar{x}(k+1) \\ e(k+1) \\ \bar{e}(k+1) \\ \hline y(k) \end{bmatrix} = \left[\begin{array}{cccc|c} F-GD & -GC & GD & GC & G_r \\ B & A & -B & 0 & 0 \\ 0 & 0 & F-LH & 0 & G_r \\ 0 & 0 & 0 & A-\bar{L}C & 0 \\ \hline H & 0 & 0 & 0 & 0 \end{array}\right] \begin{bmatrix} x(k) \\ \bar{x}(k) \\ e(k) \\ \bar{e}(k) \\ \hline r(k) \end{bmatrix}.$$

where the control law is $u(k) = C\hat{\bar{x}}(k) + D\hat{x}(k)$, and the error variables are defined as $e(k) = x(k) - \hat{x}(k)$ and $\bar{e}(k) = \bar{x}(k) - \hat{\bar{x}}(k)$. This is particular case of the configuration in sec. 2.1; the system of Fig. 5 is simply a reduced-order version of this.

In both the reduced-order and full-order versions of the system dynamics, there is coupling between $r(k)$ and $e(k)$: a manoeuvring target will introduce errors into the loop that are independent of the tracking controller. This decoupling of the stability properties does not mean that the observer cannot introduce unwanted transients into the loop if it is badly designed.

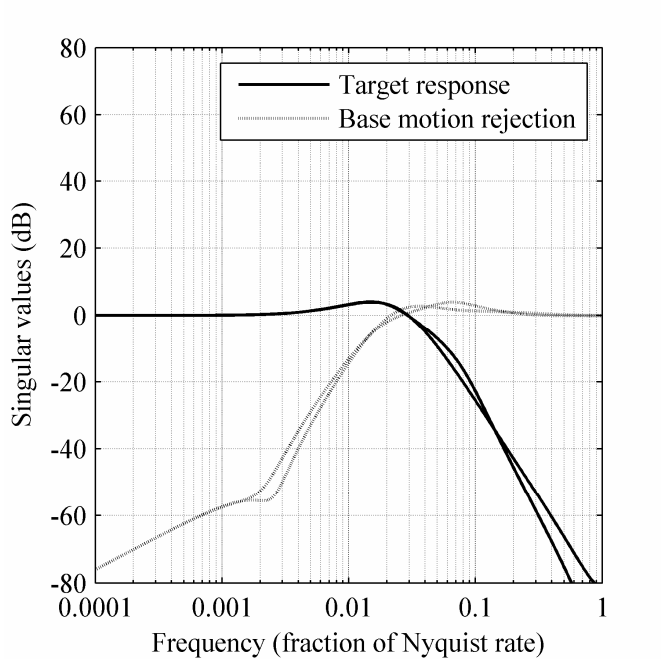

Fig. 8. Closed-loop singular values (frequency response) achieved through H_∞ loop-shaping.

3.3 Robust Control Design

Fig. 6 shows a block diagram representation of the control loop, as used for steering controller design. Note the presence of an 'inner velocity loop' present within the tracker: this is part of the legacy hardware, and will be treated as part of the plant. Feedback through the 'position' component of the pointing error returned by the observer and through angular rate measurements reported by fibre-optic gyroscopes. The nominal tracker and gyroscope dynamics are considered 'known', having been experimentally measured. All signals have two components, one in training, the other in elevation. The representation is of course a simplification: software and electronic limits are present at the input to the inner velocity loop, the gearboxes have a small backlash region and there is certainly some degree of friction in the real mount. There is also some variation in the plant's inertial characteristics according to the angular configuration of the tracker.

The H_∞ loop-shaping technique was applied using the inputs to the inner velocity loop as the design inputs, and the position error as the output; gyroscope data was available, but assigned zero-weighting because its use proved unnecessary. The position error weighting was chosen to ensure zero steady-state error in response to a target with a constant velocity; this would also provide a bounded steady-state error when tracking a target with a constant acceleration. (These can be determined using the well-known Final Value Theorem.) A normalized coprime factor H_∞ optimization was then applied. The plant singular values before and after weighting and after robustification are shown in Fig. 7. It can be seen that the robustification process has slightly "flattened off" the frequency response close to the crossover frequency. The optimization achieved $b(P_W, C_\infty) = 0.31$, which was taken to indicate good robustness.

3.4 Outline of Practical Results and Simulations

The reader will appreciate that detailed results are commercial sensitive, and it is neither possible to give precise numerical models nor detailed representations of test results.

When tested on a ground-based radar tracker observing real targets, the controller performed well with no discernable steady-state lag. A few irregularities were observed at close ranges, though these were traced to problems with the sensor and data fusion integration, and subsequently resolved. A small limit cycle was present, as expected from earlier programmes. When an artificial vibration was introduced to the loop (representing the engine-induced vibration present in a deployed system) the limit cycle disappeared. When deployed on a second tracker, no significant differences were reported, increasing our confidence that the controller was sufficiently robust to handle inter-plant variations. The control algorithm has continued in service, and simulation results predict satisfactory performance in the presence of large low-frequency platform motion.

3.5 Bounds on Controller Performance

Using the ν-gap stability bound (section 2.3), it is to be expected that the tracking loop system will be stable for all systems P_i satisfying

$$\arcsin \delta_\nu(P_W, W_o P_i W_i) < \arcsin b(P_W, C_\infty).$$

Whether this constitutes acceptable robustness is not clear: one criterion in Vinnicombe (2000) for a successful choice of weighting is that the weights should be such that expected variations between the weighted nominal plant and the weighted systems under control are small in a ν-gap sense. A possible avenue for exploring this further could be the application of formal validation techniques based on the ν-gap (Davis, 1996; Steele and Vinnicombe, 2001): these have been successfully applied in the field of flight control (Auger, 2005) and they may well be effective in assessing through-lifecycle plant variation for this scenario. An alternative area for exploration could be the nonlinear aspect of system behaviour, particularly the predictability of the small limit cycle present in the closed-loop system.

4. CONCLUSION

This paper has summarized results concerning observers and robust control from the literature (sec. 2) and applied these to a real-world marine radar tracker (sec. 3). It has been shown that for the configuration given, it is possible to separate the target tracking and sightline steering problems: the closed-loop pole locations are independent, though observation errors will still have an impact on the controller output. By considering the problem as one of reduced-order observer design, it was shown that this consistent with standard state-

feedback separation theory. Finally, H_∞ loop-shaping was used to design a steering control law: traditional frequency-response compensation techniques were used to achieve required closed-loop properties, followed by H_∞ optimization to 'robustify' the closed-loop. The ν-gap metric was used to give a deterministic bound on the set of plants that could be stabilized by the controller.

ACKNOWLEDGMENT

The authors thank the Maritime Gunnery and Missile Systems Integrated Project Team at the U.K. M.o.D. and BAE Systems Integrated System Technologies for permission to publish this work.

REFERENCES

Auger, D. J. (2005). *Modelling for Flight Control*, Ph.D. dissertation, University of Cambridge.

Balas, G., R. Chiang, A. Packard and M. Safonov (2007). *Robust Control Toolbox 3 User's Guide*. The MathWorks, Inc., Natick, MA.

Davis, R. A. (1996). *Model Validation for Robust Control*, Ph.D. dissertation, University of Cambridge.

Franklin, G. F., J. D. Powell and A. Emami-Naeini (1994). *Feedback Control of Dynamic* Systems (3rd ed.). Addison-Wesley, Reading, MA.

Franklin, G. F., J. D. Powell and M. L. Workman (1998). *Digital Control of Dynamic Systems* (3rd ed.). Addison-Wesley. Menlo Park, CA.

Glad, T. and L. Ljung (2000). *Control Theory: Multivariable and Nonlinear Methods*. Taylor & Francis, London.

Hyde, R. A., K. Glover and G. T. Shanks (1995). VSTOL First Flight of an H_∞ Control Law, *Computing and Control Engineering Journal*, **6:**11–16.

McFarlane, D. C. and K. Glover (1992). A loop shaping design procedure using H_∞ synthesis, *IEEE Trans. Automat. Contr.*, **37**:6: 759–769.

Steele, J. H. and G. Vinnicombe (2001). Closed-loop time-domain model validation in the ν-gap metric, In: *Proceedings of the 40th IEEE Conference on Decision and Control*, pp. 4332–4337.

Vinnicombe, G. (2000). *Uncertainty and Feedback: H_∞ Loop-Shaping and the ν-Gap Metric*. Imperial College Press, London.

Welch, G. and G. Bishop (1995). *An Introduction to the Kalman Filter*, TR 95-041. University of North Carolina at Chapel Hill.

Zhou, K., J. C. Doyle and K. Glover (1996). *Robust and Optimal Control*. Prentice Hall, New Jersey.