1

Introduction

1.1 Overview

1.1.1 Motivation

Historically, feedback has been used in control system engineering as a means for satisfying design constraints requiring
(i) stabilization of insufficiently stable systems;
(ii) reduction of system response to noise;
(iii) realization of specified transient-response and/or frequency-response characteristics (e.g., prescribed poles and zeroes); and
(iv) improvement of a system's robustness against variations in open-loop dynamics (e.g., parameter variations, unmodeled dynamics or non-linearity, singular perturbations, etc.).
Classical feedback synthesis techniques (viz., the graphical techniques involving Nyquist loci, Bode plots, Nichols charts, etc.) include procedures that ensure directly that each of these types of design constraints is satisfied [16, 43]. Unfortunately, the direct methods of classical feedback theory become overwhelmingly complicated for all but the simplest feedback configurations; in particular, the classical theory cannot cope simply and effectively with multiloop feedback.

Modern feedback design techniques—these include, for example, pole-placement [18], linear-quadratic-Gaussian optimal feedback [9, 10], etc.—have made relatively simple the solution of many *multiloop* control synthesis problems. The modern techniques can be readily applied in a computer-aided-design environment to provide an effective method for solving feedback design problems having constraints of the first three of the aforementioned types. However, the modern methods do not lend themselves naturally to problems in which there are design constraints of the fourth type, i.e., specifications calling for a robust tolerance of bounded uncertainty, nonlinearity, or variations in open-loop dynamics.

The inability of modern feedback synthesis techniques to handle such robustness specifications as easily and naturally as classical techniques stems from a fundamental difference between the classical and modern approaches to feedback design. In classical feedback synthesis techniques, the design model of the plant (typically a Nyquist locus) serves to specify directly the set of stabilizing scalar feedback gains for the plant model: a scalar feedback gain k is stabilizing for the plant model if and only if 1/k lies (on the real axis) in the complement of the region of the complex plane enclosed by the Nyquist locus of the plant. If the plant model is known to be accurate only to within certain bounds, then one can model this bounded uncertainty with a "fuzzy" Nyquist locus, for which the set of stabilizing feedback gains is still given by the complement of the region

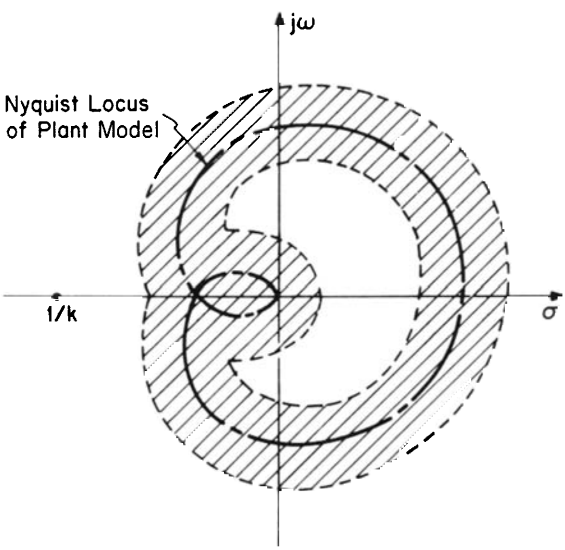

Figure 1.1 The feedback gain k stabilizes every plant whose Nyquist locus lies inside the shaded "fuzzy" band about the plant model.

enclosed by the Nyquist locus (figure 1.1). The property of classical feedback techniques that enables engineers to deal easily with robustness specifications—a property not possessed by modern multivariable feedback synthesis techniques—is this natural characterization of sets of robustly stable feedback laws directly in terms of sets of possible plant dynamics.

In the jargon of classical feedback theory, the maximal variation in plant open-loop dynamics—e.g., the maximum "fuzziness" of the plant Nyquist locus—that can be tolerated before a particular feedback design becomes unstable defines that design's *stability margins: gain margin* and *phase margin* are quantitative measures of stability margin. Engineering experience has shown that the stability margins of a feedback design provide a useful measure of its robustness against the effects of bounded variation in open-loop dynamics; consequently it is common for feedback controller design specifications to include requirements for prescribed minimal gain margin and phase margin.

The fundamental objective of the research reported here has been to develop a means for incorporating robustness specifications—i.e., design specifications calling for prescribed stability margins—into modern multivariable feedback design procedures, including procedures employed in multivariable stochastic model-reference estimator design.[1] One may appropriately view this fundamental objective as being the development of a *theory of approximations* applicable to multivariable dynamical systems, a theory that quantifies the trade-offs between modeling approximations and feedback-law choice by associating with each feedback law stability margins establishing limits to the tolerable imprecision in the system model. Such modeling approximations arise routinely in every engineering problem, not only as a consequence of the unavoidable uncertainly associated with physical processes, but also as a consequence of intentional model simplification. The latter includes such common practices as linearization, neglect of weak coupling between subsystems in decentralized

1 Recall that a *model-reference* estimator is an estimator incorporating an internal model of the process dynamics. In such estimators "residual error" (which is the difference between the model output and the observed process output) is fed back to the internal model so as to control the estimate error [104, p. 403]. Virtually all practical recursive estimator designs (including for example the Kalman filter, the extended Kalman filter, and the Luenberger observer) are model-reference estimators.

designs, and time-scale decomposition into low-, medium-, and high-frequency system models; the latter is a standard engineering strategy employed in developing simplified hierarchical designs.

Accomplishing this fundamental objective has entailed the development of

(i) a means for specifying stability margins for multiloop feedback systems (since classical single-loop characterizations of stability margins, such as gain and phase margin, are in general inadequate for characterizing multiloop stability margins); and

(ii) a new stability theory that provides a direct characterization of sets of robustly stable multiloop feedback laws in terms of "fuzzy" sets of possible multiloop plant dynamics.

1.1.2 Description of Results

The central result described in this book—a fundamental theoretical result which forms the foundation of the entire work—is an abstract and extremely powerful new stability theorem (theorem 2.1). In essence, the theorem shows that a multiloop feedback system is closed-loop stable if there exists a topological separation (into two disjoint sets) of the function space on which the system's dynamical relations are defined, the relation for the forward loop lying in one part of the separated space and the relation for the feedback loop lying in the other. This theorem has a direct bearing on the robustly stable feedback synthesis problem: if one part of the separation is taken to be a bounded region about the plant model and the other part is taken to be the complement of this region, then the theorem provides a direct characterization of the set of robustly stable feedback laws as the complement of the set of possible plant input-output relations, just as the Nyquist theorem does for single-loop classical feedback designs having a nondynamical scalar feedback gain k. The input-output relations may be specified by, for example, transfer function matrices or, possibly, nonlinear state equations: in the former case the topological separation condition can be verified in the frequency domain, leading directly to powerful multiloop generalizations of the circle and Popov stability criteria (see section 2.5); in the latter case one finds that the classical Lyapunov stability theory emerges as a special case in which a positive-definite Lyapunov function is used to establish the topological separation (corollaries 2.1 a, b). Inasmuch as the new stability theorem includes as special cases some of the most powerful existing stability theorems (e.g., Lyapunov, Popov, or circle stability criteria) and inasmuch as it provides a fundamentally new perspective for stability theory (viz., topological separation of function spaces), it constitutes a new theory of stability—the theorem has been designated the *main stability theorem.*

A methodology based on this main stability theorem has been devised for multiloop feedback system robustness and stability margin analysis (section 2.6). The methodology, which is in much the same spirit as the classical approach to single-loop robustness problems, uses frequency-dependent "sector conditions" to characterize the frequency-dependent "fuzziness" in a *multivariable* design model. A special case of the main stability theorem called the *sector stability theorem* (theorem 2.2) together with various *frequency-domain* tests (section 2.5) of the conditions of the sector stability theorem provide a practicable means for specifying multivariable feedback system stability margins and a theoretical basis for the design of robust multiloop feedback laws to meet such specifications; a conceptual computer-aided design procedure is outlined in section 2.6. Potential applications of the results include simplified hierarchical control for large systems and gain scheduling for adjustable set-point nonlinear regulator systems.

The implications of the new theory with regard to the stability margins of modern multivariable linear-quadratic-Gaussian (LQG) optimal estimators and controllers are examined in detail. The continuous-time case is considered in chapter 3, and the discrete-time/sampled-data case in chapter 4.

The design-specific stability margins of optimal linear-quadratic state-feedback (LQSF) designs are characterized in terms of the system matrices, quadratic-performance-index weighting matrices, and the optimal solution of the Riccati equation: these margins are characterized as a convex set of nonlinear dynamical deviations between the design model and the actual plant that can be tolerated without inducing instability. Additionally, it is shown that the continuous-time LQSF optimal design procedure is inherently robust in that it automatically ensures certain minimal stability margins including an infinite gain margin, at least a $\pm 60°$ phase margin, and at least a 50% gain reduction tolerance at each control input channel; discrete-time/sampled-data designs are found to approximate this inherent robustness, but the robustness is degraded as the sampling interval increases.

Viewing the Kalman filter as a feedback system in which residual error is fed back to control the estimation error, the duality between Kalman filters and LQSF controllers is expoloited to provide a characterization of the nonlinearity tolerance of a constant-gain extended Kalman filter (CGEKF); the CGEKF is a nonlinear extended Kalman filter employing a precomputed constant residual-gain matrix and having drastically reduced on-line computational requirements relative to extended Kalman filters employing a time-varying residual-gain matrix that is adaptively updated on-line. The results include analytically verifiable conditions that can be used to confirm the (global!) stability of CGEKF designs and, provided that the CGEKF incorporates an accurate internal model of the nonlinear process dynamics, ensure the nondivergence of CGEKF estimates. The CGEKF design procedure is found to have an inherent robustness which can be interpreted as including an infinite gain margin, at least $\pm 60\%$ phase margin, and at least 50% gain reduction tolerance at each sensor output channel; though it should be noted that, because an accurate internal system model is required for CGEKF nondivergence, the CGEKF robustness results do not have the same interpretation as the dual LQSF robustness results.

The CGEKF and LQSF robustness results are shown to combine in a fashion reminiscent of the separation theorem of estimation and control to suggest a powerful technique, based on linear-quadratic-Gaussian optimal feedback theory, for the synthesis of simplified dynamical output-feedback compensators for nonlinear regulator systems. The technique leads to a feedback compensator design consisting of a cascade of a CGEKF and an optimal constant LQSF gain matrix. It is proved that the inherent robustness of optimal linear-quadratic state feedback against unmodeled nonlinearity combines with the intrinsic robustness of the CGEKF to assure that such feedback designs will be closed-loop stable even in systems with substantial nonlinearity, assuming that the CGEKF incorporates an accurate internal model of the nonlinear plant dynamics.

The aforementioned LQSF regulator stability results and CGEKF stability and nondivergence results are derived in the general context of the class of constant-gain controllers and nonlinear estimators whose design is not necessarily based on statistical considerations; for example, nonlinear estimator designs intended to optimize structural simplicity or error-transient response, i.e., nonlinear observers [107]. This general class of constant-gain controllers and nonlinear estimators includes as special cases LQSF and CGEKF designs. In the context of this broader class of suboptimal nonlinear controllers and estimators, the results provide analytically verifiable conditions which can be used to test for nondivergence and stability and to evaluate robustness against the effects of design approximations; though one cannot in general expect such designs to be as robust as LQSF and CGEKF designs. The output-feedback separation-type property applies to this broader class of controllers and estimators, showing that nondivergent estimates can, unconditionally, be substituted for true values in otherwise-stable feedback systems without ever causing instability.

It is demonstrated that all the CGEKF/LQSF nondivergence and stability results extend to state-augmented designs (section 3.8 and 4.8). Such state-augmented designs include, for example, proportional-integral controller designs which track with zero steady-state error and compensated CGEKF nonlinear estimator designs that have zero steady-state bias error.

1.1.3 Structure of Monograph

This monograph consists of five chapters:

1 Introduction

2 Stability and Robustness: A Geometric Perspective and Frequency-

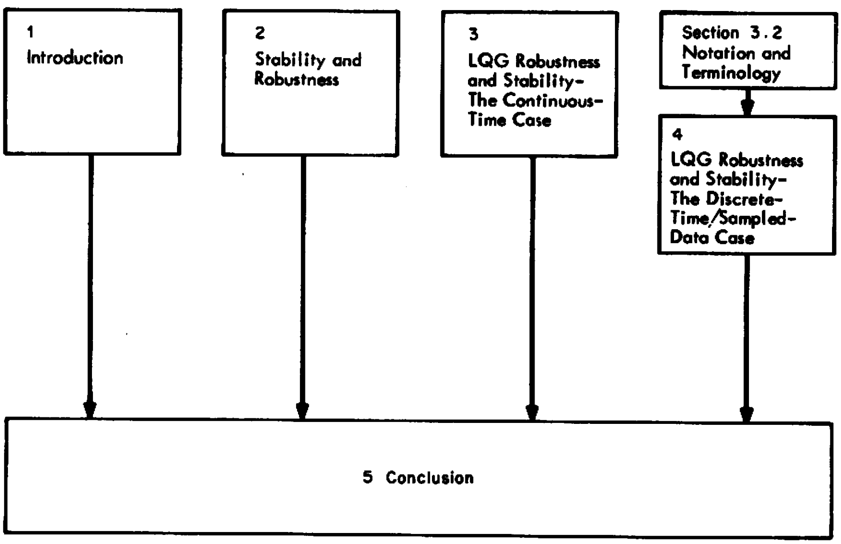

Figure 1.2 Logical relationship of chapters.

Domain Criteria
3, 4 LQG Robustness and Stability
5 Conclusion.

The logical interdependence of the various chapters is illustrated in figure 1.2. The monograph has been arranged so that the first four chapters are logically independent and may be read in any order, though it is recommended that the sections within each chapter be read in sequence. The proofs of the results in chapters 3 and 4, all of which are based on the results developed in chapter 2, are contained in the appendix, a few exceptions being made for proofs that are especially simple or illuminating. The developments in chapters 3 and 4, which concern respectively the continuous- and discrete-time LQG problem, are completely parallel, even to the numbering of sections and theorems; readers wishing to compare the results developed for the two cases will find this parallelism helpful.

1.2 Previous Work and Related Literature

1.2.1 Robustness and Sensitivity of General Feedback Systems

The fundamental work on the robustness of feedback systems is due to Bode [16, pp. 451–488]. Bode's work concerned both the robustness of feedback systems and the closely related issue of differential sensitivity, although Bode did not use these terms. The term *robustness* refers to tolerance of large disturbances (lying within specified bounds) whereas the term *differential sensitivity* concerns the effects of vanishingly small disturbances.

Bode showed that the effects of vanishingly small perturbations in the gains of a single-loop linear feedback system are directly related to the return ratio (or loop gain) and return difference (one minus the return ratio) of a feedback system. Bode's results on differential sensitivity have since been extensively studied [43] and extended [24, 53] to multiloop, nonlinear, time-varying feedback systems. A discrete-time notion of differential sensitivity has also been formulated by Kwakernaak and Sivan [55, p. 427]. It appears that the analysis of differential sensitivity is now well understood; though the development of methods for synthesizing feedback systems with reduced differential sensitivity continues to be an area of active research [37]. Reference [26] is an anthology of many of the key journal articles which have appeared on the subject of differential sensitivity. Reference [25] also contains several relevant articles.

In the area of robustness, Bode's principal contribution was the observation that the amount of tolerable uncertainty in the open-loop dynamics of a single-loop linear feedback system can be expressed in terms of a region in a space whose "points" are open-loop system transfer functions, each such "point" being described by the Nyquist locus of its associated transfer function. Specifically, Bode showed how the notions of gain and phase margin can be exploited to arrive at a simple and useful means for characterizing these regions of tolerable uncertainty. The engineering implications are developed in detail by, e.g., Horowitz [43]. It is shown in [44] that classical Nyquist-Bode theory can be iteratively applied in order to design robustly stable multiloop feedback having the special structure of "cascaded multiloop feedback." Rosenbrock [88; 89, pp. 198–208; 90], McMorran [66], and Belletrutti and MacFarlane [15, 64, 65] have developed and examined the design implications of a multiloop generalization of the Nyquist stability criterion. However, except for the narrow class of "diagonally dominant" (i.e., "weakly-coupled") systems [20, 90, 92], the quantitative robustness implications of the multiloop Nyquist results remain unclear.

Results having a close relation to the issue of feedback system robustness have been developed by Zames [118, 119]. Viewing feedback systems in terms of the dynamical input-output relations of their components, Zames showed that "conic sectors" in an appropriate space of input-output relations can be used to aggregately characterize complex nonlinear dynamical input-output relations in a simple, useful fashion. While the primary emphasis in Zames's work was the use of conic sectors to provide a simple characterization of complex nonlinear relations, he made the following brief observation in his conclusions [118]:

> One of the broader implications of the theory here concerns the use of functional analysis for the study of poorly defined systems. It seems possible, from only coarse information about a system, and perhaps even without knowing details of internal structure, to make useful assessments of qualitative behavior.

Despite the appearance of myriads[2] of publications expanding upon and refining the results in Zames's paper [118], the idea of using functional analysis for poorly defined systems has not previously been fully developed. Nevertheless, this idea is implicit in the classical frequency-domain ap-

2 Among these myriads are, to name a few, [6], [22], [23], [27], [32], [42], [56], [67], [86], [106], [110], and [119]; a complete listing would probably include at least several hundreds of references.

proaches to robust feedback design[3] and this idea forms the thesis upon which the multiloop feedback and stability margin results of this book are built. In particular, it is Zames's work [118, 119] which laid the foundation and provided much of the inspiration for the theory of stability developed in chapter 2.

It should be noted that some authors (e.g., Davison [29]) have used the term "robustness" in connection with linear feedback systems that track with zero steady-state error in the presence of additive disturbance inputs satisfying specified dynamical equations. It is demonstrated in [29] and [30] that if an appropriate servo-compensator is employed then such zero steady-state error tracking occurs and is robust against plant modeling errors, provided that the overall system remains closed-loop stable. In the present monograph, we dwell principally on a different, but related, aspect of robustness: we are concerned primarily with the amount of plant modeling error that can be tolerated before a feedback system becomes unstable.

1.2.2 Linear-Quadratic-Gaussian (LQG) Estimators and Feedback Controllers—General

One of the more powerful approaches to multivariable feedback system design, an approach whose robustness implications are examined in detail in this monograph, is the linear-quadratic-Gaussian (LQG) procedure. The technical and philosophical issues relating to the application of the LQG procedure are discussed by, for example, Athans [9, 10]; the textbooks by Anderson and Moore [5] and Kwakernaak and Sivan [55] are excellent sources of detailed information about the LQG procedure.

Briefly, the LQG method is an algorithm for synthesizing output-feedback compensators which minimizes a designer-selected quadratic performance index for a linear plant subjected to Gaussian white noise of known mean and covariance. Based on the *separation theorem* of estimation and control, an LQG compensator can be split into two parts: a Kalman filter, which estimates the state of the plant, and a memoryless state-feedback gain, which acts on the state estimate generated by the Kalman filter to generate compensating inputs to the plant. In the special case in which the plant output includes noise-free measurements of the

3 Frequency-domain methods and transfer function methods are in fact function space methods. For example, a transfer function may be viewed naturally as a functional operator mapping complex input functions (defined over the complex plane) into complex output functions (likewise defined over the complex plane).

plant state, the Kalman filter is eliminated and the compensator is called a *linear-quadratic state-feedback* (LQSF) regulator.

The LQG technique provides a straightforward means for synthesizing stable multiloop linear feedback systems which are insensitive to Gaussian white noise. Variations of the LQG technique have been devised for the synthesis of feedback systems with specified poles and eigenvectors [5, pp. 77–78; 45; 72; 115; 126] and for systems with constraints on controller structure [19, 59, 60, 101]. By state-augmentation methods the LQG technique can be applied to the design of feedback compensators for systems subject to persistent disturbance inputs and to other types of nonwhite noise [17, 47–49, 57, 101, 102, 117]. It is well known [8, 9] that by substituting a nonlinear extended Kalman filter for the optimal linear estimator (i.e., the Kalman filter) in the LQG controller, one may adapt the LQG procedure to suboptimal nonlinear control problems.[4] Thus, LQG theory has come to play a central role in much of modern multiloop feedback theory.

1.2.3 Stability Margins, Sensitivity, and Robustness of LQG Estimators and Controllers

The fundamental work on the differential sensitivity properties associated with LQG systems is due to Kalman [52]. Kalman showed that (under mild assumptions) a linear single-input state-feedback system is optimal with respect to some quadratic performance index if, and only if, the system is stable and has a return difference of magnitude greater than unity at all frequences. Kalman noted that classical control theory requires this condition on the return difference for reduced sensitivity to component variations. Kalman's sensitivity results have since been more precisely interpreted by Perkins and Cruz [82] and generalized to multi-input LQSF regulators by Anderson [1] and MacFarlane [63]. Perkins and Cruz [83] summarize these results, and Wonham [120, ch. 13] contains a brief critical discussion. The differential sensitivity of the optimal cost to parameter variations is addressed by Barnett [11, 13, 14]. The effects on the optimal quadratic cost of variations in sampling interval is addressed by Levis et al. [61]. Techniques for synthesizing LQSF designs with reduced sensi-

4 While this nonlinear extension of LQG feedback theory is well known, its rigorous justification rests on the nonlinear separation theorems developed in chapters 3 and 4 of this monograph, viz., theorems 3.3 and 4.3.

tivity are proposed in [51, 103].

Relatively little has been produced regarding the differential sensitivity properties associated with Kalman filters and with general LQG regulators incorporating Kalman filters in their feedback loops. With only limited success, Anderson [2] attempted to determine under what circumstances a general LQG regulator is equivalent to an LQSF regulator. Kwakernaak [54] showed (under the assumption that the controlled plant is minimum phase) that in the limit as the control cost-weighting matrix goes to zero (i.e., as the loop gain increases so that the poles of the Kalman filter dominate the system response), general LQG regulators do exhibit reduced differential sensitivity; however, it is not clear that an LQG design dominated by the poles of its filter would be satisfactory in other respects.

Regarding the robustness properties of LQSF systems, perhaps the most significant result is due to Anderson and Moore [5, pp. 70–76]. Employing the Nyquist and the circle stability theorems, Anderson and Moore showed that Kalman's result [52] concerning the return difference of LQSF systems can be used to conclude that single-input LQSF regulators have $\pm 60^\circ$ phase margin, infinite gain margin, and 50% gain reduction tolerance. Moreover, they showed that the gain properties apply to nonlinear time-varying gains. Related results by Moylan [74, 75] generalize results of Barnett and Storey [12] and of Moore [73] to parameterize classes of memoryless, nondynamical feedback perturbations which will not destabilize a multi-input optimal state-feedback regulator with quadratic performance index. Less general multiloop results have been derived independently by Wong [113, 114]. Gilman and Rhodes [36] have developed an upper bound on the quadratic performance index of suboptimal nonlinear LQSF designs; provided the system is cost-observable and provided the bound does not "blow up," their results can be used to evaluate the gain margins and nonlinearity tolerances of LQSF designs. Recent papers dealing with LQSF robustness include [80], [81], and [108], which are in part based on some of the ideas that have been developed in the course of the research reported here. The papers [95]–[97] contain preliminary reports of the result (section 3.4) that all continuous-time LQSF designs have infinite gain margin, at least $\pm 60^\circ$ phase margin, and at least 50% gain reduction tolerance at each control input channel; though it should be noted that the methods of proof employed in [95]–[97] are different.

As an indication of at least one of the limitations of the robustness properties universal to LQSF regulators, Rosenbrock and McMorran

[91] show by means of an example that LQSF regulator designs may be conditionally stable. That is, the failure of a single loop in a multiloop LQSF regulator may destabilize a system which is open-loop stable. Wong [113] examines this problem in greater depth.

The literature on the subject of robustness and computational considerations in nonlinear estimation is sparse and largely inconclusive. The discussion of nonlinear estimation in Schweppe [104, ch. 13] provides a good intuitive understanding of the trade-offs between computational requirements and residual-gain choice; though the possibility of a constant residual gain is not explicitly considered. The idea of using a constant residual gain for linear filtering is well known [34, pp. 238–242], but the connection with nonlinear filtering has not been established. Of the existing literature on nonlinear estimation, [35] and [107] appear to be the most closely related to the present work.

Gilman and Rhodes [35] suggest a procedure for synthesising nonlinear estimators with a pre-computable but time-varying residual gain. Their estimator, like the extended Kalman filter, has the intuitively appealing structure of a *model-reference* estimator [105, p. 403]; i.e., it consists of an internal model of the system dynamics with observations entering via a gain acting on the residual error between the system and model outputs. The distinguishing feature of the estimator suggested in [35] is that the residual gain is chosen so as to minimize a certain upper bound on the mean-square error. This procedure tends to ensure a robust design since, assuming the minimal value of the error bound does not "blow up," the estimator cannot diverge. A limitation of this design procedure is that the error bound may be very loose for systems with substantial nonlinearity; so there is no assurance that the bound-minimizing residual gain is a good choice. Also, there is no a priori guarantee that the resultant estimator will even be stable since the minimal error bound may become arbitrarily large as time elapses. Similar results are developed for discrete-time systems by Gusak and Simkin [38].

Tarn and Rasis [107] have proposed a constant-gain model-reference nonlinear estimator which is a natural extension of Luenberger's observer for linear systems, having a design based solely on stability considerations. The results of [107] show that, given such a nonlinear observer design, if certain Lyapunov functions can be found, then one can conclude that

(i) The estimator is nondivergent; and

(ii) The estimator can be used for state reconstruction in a full-state-feedback system without causing instability.

However, from an engineering standpoint the results of [107] are deficient in that they are nonconstructive: no design synthesis procedure is suggested; no method is proposed for construction the Lyapunov functions required to test the stability of a design; no procedure is suggested for minimizing the estimator's error. The CGEKF results presented in this monograph (sections 3.5–3.7 and 4.5–4.7) address all these deficiencies by providing a constructive procedure for synthesizing stable constant-gain model-reference estimator designs which are to a first approximation optimally accurate. Moreover, these results prove that, provided the estimator is nondivergent, it can be used for state reconstruction without ever causing instability, irrespective of the availability of Lyapunov functions.

The papers of Patel and Toda [121, 122], in contrast to the aforementioned papers on nonlinear estimation and in contrast to the present work, do not require an exact internal model of the process dynamics in their results regarding the robustness of optimal estimators. These authors have developed an upper bound on the mean-square error of mismatched optimal linear estimators, i.e., mismatched Kalman filters. However, their results have the important limitation that they apply only when the control input to the process is identically zero.

Results concerning the robustness of general LQG feedback controllers have been reported by Gilman and Rhodes [35]. These authors have developed an upper bound on the expected value of the quadratic performance index for white-noise-driven, continuous-time, output-feedback systems employing a state-feedback gain matrix acting on a state estimate generated by a nonlinear model-reference estimator which, like the extended Kalman filter, incorporates a nonlinear internal model of the plant dynamics. Provided the system is "cost-observable," provided the upper bound on the cost does not grow unboundedly as time elapses, and provided the nonlinear model incorporated in the estimator exactly models the plant dynamics, then this result can be used to assure the closed-loop stability of the overall system. Since in practical applications the actual value assumed by the performance index is often not of particular engineering interest and since the upper bound on the cost in [35] may be quite loose for systems with substantial nonlinearity, the principal engineering interest in these results appears to be with regard to their stability implication. On the other hand, the separation theorem for nonlinear systems developed in this monograph (theorem 3.3) allows one to ascertain the stability of such systems much more simply and directly.

Underlying the various conclusions about the minimal *inherent* robustness[5] and *inherent* insensitivity of LQG estimator and controller designs are certain more basic properties associated with optimality. The problem of characterizing the various properties associated with optimality is intimately related with the so-called "inverse problem" of optimal control, which concerns the characterization of the set of performance indices for which a given control law is optimal. Related papers are [1], [52], [71], [75], and [116].

5 For example, the inherent infinite gain margin and $\pm 60°$ phase margin robustness properties of single-input LQSF systems proved by Anderson and Moore [5, pp. 70–76].