

**MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LABORATORY FOR INFORMATION AND DECISION SYSTEMS
CAMBRIDGE, MA 02139**

FINAL REPORT

ON

ROBUST STOCHASTIC ADAPTIVE CONTROL

**ONR CONTRACT N00014-82-K-0582
NR 606-003
MIT OSP NO. 92775**

PREPARED BY:

**PROF. LENA VALAVANI
PROF. MICHAEL ATHANS**

LIDS - FR - 1744

JANUARY 28, 1988

SUBMITTED TO:

**DR. JOHN R. CANNON
CODE 1111
MATHEMATICAL SCIENCES DIVISION
OFFICE OF NAVAL RESEARCH
ARLINGTON, VA 22217-5000**

0. SUMMARY

In this final report we summarize the activities of the MIT/LIDS research group for the time period 1 September 1982 to 31 December 1987. The research, funded by ONR contract N00014-82-K-0582, deals with fundamental issues in robust adaptive control systems, and the potential application of advanced control system design methodologies to the multivariable control of submarines.

The research conducted was highly successful, and had significant (and controversial) impact upon the theory of adaptive control. The research of Rohrs, Valavani, Athans and Stein pointed out potential instabilities of then existing adaptive control algorithms caused by the presence of unmeasurable output disturbances and high frequency unmodeled dynamics. The publications of Rohrs *et al* were instrumental for defining new research directions in the adaptive control field, and the topic of *Robust Adaptive Control* became a new area for worldwide research. The research of Krause *et al* provided the first direction for the use of what is now called *Averaging Theory* for the analysis of adaptive control algorithms in the presence of disturbances and unmodeled dynamics. The research of Orlicki *et al* provided the first set of adaptive algorithms that actively employ real-time signal processing to compute frequency domain parameters which can be used to safeguard the stability of Model Reference Adaptive schemes that employ *Intermittent Adaptation*. The research of LaMaire *et al* deals with novel formulation of *Hybrid Robust Identification* algorithms which identify in real-time both time-domain models of the unknown plant and modeling error bounds in the frequency domain. The research of Milich *et al* develops theory and methodologies for designing robust compensators, with guaranteed performance in the presence of large structured and unstructured plant uncertainties. This complements the research conducted which helped streamline the LQG/LTR design methodology for non-adaptive feedback systems. Finally, our research on the design of multivariable control systems for modern submersibles have brought into focus the advantages of present design methodologies and helped pinpoint directions for future theoretical research.

The funds provided by the Office of Naval Research provided whole or partial support for two faculty, five Ph.D. and four M.S. graduate students during the contract time period.

1. ADAPTIVE CONTROL THEORY

1.1 The Robust Adaptive Control Problem.

Our research support under this ONR contract started shortly following the completion of the Ph.D. thesis of C.E. Rohrs under the supervision of Professors Valavani, Athans, and Stein. This research had uncovered major shortcomings with available adaptive control algorithms, which were proven to be globally stable under certain mathematical assumptions. We showed by a combination of analysis and simulations that existing adaptive control algorithms could become unstable in the presence of unmodeled dynamics and unmeasurable disturbances. Our research, eventually documented in publications [1], [2], [8], [9], [10] and [14], was originally received with open hostility in the Decision and Control Conferences and the American Control Conferences and resulted in many heated discussions.

Eventually, by 1985, the adaptive control community became convinced that existing adaptive control algorithms could break into instability. The so-called *Rohrs et al counterexample*, fully described in [14], became the test benchmark by which modifications of adaptive algorithms were tested on. Soon a new field of international research on the *Robust Adaptive Control Problem* was born. Research on this topic is vigorously pursued by many distinguished researchers at present; nobody as yet has arrived on a simple modification to the original adaptive algorithms that preserves global stability and robustness to unmodeled dynamics.

1.2 The Beginning of Averaging Theory.

The results of Rohrs *et al* pointed out that we needed new tools for understanding the mechanisms of stability and instability in adaptive control systems. J. Krause addressed this key problem in his M.S. thesis, see publications [3] and [4], and suggested in a preliminary form a method of analysis which averaged out the slow transients of the system. This set of results were later on extended by many other researchers in adaptive control using ideas from both singular perturbation theory and nonlinear oscillation theory. This area of research is currently known as *Averaging Theory* and provides a more rigorous way for explaining the complex mechanisms that give rise to the instability phenomena reported by Rohrs *et al*.

1.3 Intermittent Adaptation and Variable Dead-Zones.

The results of Rohrs *et al* and Krause *et al* pointed out that a potential villain in the destruction of adaptive control stability was that the combination of certain types of reference inputs, disturbances, and unmodeled dynamics provided spurious, and

unwanted, information to the (explicit or implicit) adaptive identification scheme. These errors, unless accounted for, could interact with the feedback mechanism and result in instability. Hence, we decided to initiate a research effort that would desensitize the adaptive system from such spurious information. Similar philosophy was followed by other researchers, e.g. Peterson and Narendra, by the use of a fixed dead-zone whose width was adjusted *a priori* based upon estimates of the size of the unknown disturbances. Only output error signals that exceeded the dead-zone were used to update the parameters of the adaptive compensators. The problem was that this dead-zone could be very conservative; also, previous researchers did not account for the impact of high-frequency modeling errors. These unmodeled dynamics could interact with both reference inputs and disturbances and introduce additional spurious signals that would confuse the identification algorithm.

The doctoral thesis of D. Orlicki, under the supervision of Professors Valavani, Athans, and Stein addressed this class of problems. We focused upon the philosophy of *Intermittent Adaptation* realized by passing the output error through a variable dead-zone; the size of the dead-zone was varied in real time by carrying certain computations, over and above those necessary to implement the classical adaptive algorithms. In this research, documented in publications [7] and [11], we were able to develop new algorithms, of the MRAC type, which have guaranteed local stability properties in the presence of unmodeled dynamics and unmeasurable disturbances. The instability of the classical MRAC schemes was prevented by the intermittent adaptation; as discussed above, this technique prevents the updating of uncertain plant parameters whenever the identification information is of dubious quality due to the simultaneous presence of unmodeled dynamics and disturbances which cannot be measured. Thus, we only adapt whenever we are sure that the real-time signals contain relevant information.

It is a highly nontrivial matter to decide, in real-time, when to adapt and when to (temporarily) stop the adaptation. The new algorithms of Orlicki *et al* involve the real-time monitoring of easily measurable signals, and require the capability of computing discrete Fast Fourier transforms (DFFT's) for those signals. Intermittent adaptation is implemented by blending the real-time spectral information generated by the DFFT's with variants of the model reference algorithms. The algorithms can be implemented through the use of a dead-zone nonlinearity whose width changes in real time based upon the DFFT calculations. To the best of our knowledge, this is the first time that an adaptive control algorithm had been developed that requires extensive real-time spectral calculations so as to guarantee stability-robustness. Due to the very significant real-time computational requirements only limited simulation results were obtained; these results were encouraging but could not be used with confidence to pinpoint the advantages and shortcomings of this class of algorithms in a practical setting.

One can question the practical utility of adaptive algorithms that require so many spectral calculations to control a relatively simple process. Nonetheless, one should not lose sight of the experience of adaptive signal-processing in which spectral calculations to improve performance are used routinely. The adaptive control problem is much harder than the adaptive signal processing problem, because in addition to improved performance one has to worry also about the stability of the adaptive feedback control problem.

Although our intermittent adaptation algorithms represent an advance in the state of the art, and undoubtedly will become controversial because of their increased computational requirements, nonetheless the most important by-product of that research was a detailed appreciation of the immense complexity of the adaptive control problem. In point of fact, we become convinced that new and different approaches to the robust adaptive control problem must be developed. There are simply too many hard questions, only tangentially related to adaptive control, that must be posed first, and of course answered, before we can proceed with confidence to using adaptive control to regulate physical systems, and especially multivariable ones. These questions motivated our subsequent research.

1.4 Robust Adaptive Identification in the Time and Frequency Domains.

Classical adaptive control algorithms use a postulated dynamic system order, i.e. a transfer function with fixed numbers of poles and zeros, and then use (explicit or implicit) identification to improve the prior estimate of the model uncertain parameters. In robust adaptive control this is necessary, but by no means sufficient. What is required is the development of a new class of adaptive identification algorithms which, with a finite amount of data, produce not only a better nominal model, but in addition generate a bound in the frequency domain that captures the presence of possible high-frequency model errors. Such bounding of model errors in the frequency domain is required by all nonadaptive design methods so as to ensure stability-robustness by limiting the bandwidth of the closed-loop system. Such identification algorithms did not exist in the classical identification literature; such questions were not even posed. Thus, we believed that it was essential to develop such algorithms and then to incorporate them in the adaptive control problem. A major milestone along these lines has been completed with the publication of Richard LaMaire's doctoral thesis, under the supervision of Professors Valavani and Athans; see publications [18] and [19].

We view the robust adaptive control problem as a combination of a robust identifier (estimator) and a robust control-law redesign algorithm. Current robust control

design methodologies, such as the LQG/LTR methodology, require: 1) a nominal model, and 2) a frequency-domain bounding function on the modelling error associated with the nominal model. A new robust estimation technique, which we call a 'guaranteed' estimator, has been developed to provide these two pieces of information for a plant with unstructured uncertainty and an additive output disturbance. This guaranteed estimator uses parametric time-domain estimation techniques to identify a nominal model, and non-parametric frequency-domain estimation techniques to identify a frequency-domain bounding function on the modelling error. This bounding function is generated using discrete Fourier transforms (DFT's) of finite-length input/output data.

Several assumptions are required by the guaranteed estimator. In addition to a priori assumptions of the structure of the nominal model along with coarse, worst-case values of the parameters, we assume that the unmeasurable disturbance is bounded and that a magnitude bounding function on the Fourier transform of the disturbance is known. Further, we assume prior knowledge of a bounding function on the unstructured uncertainty of the plant relative to our choice of nominal model structure. These assumptions allow our time-domain estimator to be made robust to the effects of unstructured uncertainty and bounded disturbances. That is, our time-domain estimator updates the parameters of our nominal model only when there is good (uncorrupted) information. Similarly, the frequency-domain estimator, which has been developed, only updates the model and current bounding function on the modelling error when there is good information. In summary, the guaranteed estimator provides a nominal model plus a guaranteed bounding function, in the frequency-domain, as to how good the model is. Accuracy guarantees in the identifier part of the adaptive controller can be used by the control-law redesign part of the adaptive controller to ensure closed-loop stability, assuming the control-law is updated sufficiently slowly.

All the equations necessary to simulate the performance of these identification algorithms were coded and debugged. Because of the extensive real-time spectral calculations, we decided to use the CYBER supercomputer at Princeton which is available for use by the MIT community at no cost for CPU time. Numerical examples which are simple enough to demonstrate the ideas yet rich enough to capture the potential pitfalls have been designed and simulated.

The simulation results indicate that for the systems tested the time-domain identification algorithm did not work very well. On the other hand, the frequency-domain algorithms worked much better.

In closed-loop identification simulations the richness of the command signal was often not sufficient to excite the plant dynamics so that the identification algorithms

could work properly. For this reason, we developed an "intelligent" scheme which would monitor the progress of the identification algorithm and inject probing signals at the appropriate frequencies at the plant input so as to enhance identification. Of course, this would deteriorate (temporarily) performance since a disturbance was injected intentionally in the feedback loop. Better identification, accompanied by higher loop-gains and bandwidths, would improve overall command-following and disturbance-rejection performance after the probing signals were terminated.

The algorithms require extensive real time computations. For sluggish plants the computational requirements are not severe. However, in order to identify and control plants with very lightly damped dynamics truly extensive CPU requirements exist. For example, in our simulation studies involving a second order plant with lightly damped poles the Cyber 205 supercomputer was too slow, for real time control, by a factor of two so as to achieve a closed-loop bandwidth of 5 rad/sec.

These findings cast a tone of pessimism, with respect to CPU requirements, in using real-time identification and high-performance adaptive control for typical aerospace plants that are characterized by lightly damped dynamics and dominant high-frequency modeling errors. On the other hand, parallel computer architectures can be exploited in this class of algorithms. Thus, more research along these lines is required.

1.5 Best Nonadaptive Compensator Design for Performance-Robustness.

Our research to date has pinpointed the need for a good initial guess for an adaptive compensator, whose parameters are then updated by the adaptive algorithm. We are developing techniques that design the best (from the viewpoint of good command-following and disturbance-rejection) nonadaptive compensator for the given prior plant uncertainty information.

In his doctoral research Mr. David Milich, under the supervision of Professors Athans and Valavani, has developed a design technique which will yield the "best" fixed-parameter nonadaptive compensator for a plant characterized by significant unstructured uncertainty; see publications [17] and [20]. The "best" compensator is defined as the one that meets the posed performance (i.e. command-following, disturbance-rejection, insensitivity to sensor noise) specifications and stability-robustness over the entire range of possible plants.

Such a robust design technique will prove useful in a number of ways. First, it will

yield a systematic procedure for designing feedback systems for uncertain plants with both stability and performance guarantees, not only for the nominal plant but for the entire set of uncertain plants considered. Thus, the feedback loop will be guaranteed to be stable and, in addition, will meet minimum performance specifications for all possible plant perturbations. Second, the solution of this robust design problem will also enable us to quantitatively address one of the most fundamental questions in adaptive control: *what are the performance benefits of adaptive control?* While much attention has been paid to the development of many specific adaptive algorithms, very little consideration has been given to this issue which is, we believe, at the heart of the adaptive control problem. Practical adaptive systems rely upon external persistently exciting signals (to ensure good identification), slow sampling (which helps stability-robustness to unmodeled high frequency dynamics) in addition to extensive real-time computation (to provide safety nets and turn-off the adaptive algorithm when it exhibits instability). All these "fixes" degrade command-following and disturbance-rejection performance and tend to neutralize the hoped-for benefits of an adaptive compensator. In light of these circumstances it is imperative that the decision to use adaptive control, for a real engineering application, must be based upon a quantitative assessment of costs and benefits.

Some of the key issues, and severe difficulties, in the design process have been identified. Conditions for stability-robustness and performance-robustness in the presence of significant unstructured uncertainty have been developed. An a-priori magnitude bound, as a function of frequency, on the unstructured uncertainty is assumed known. In order to reduce the conservatism of the stability and performance conditions with respect to the structured uncertainty, directional information (in the complex plane) associated with the plant-parameter variations is exploited. Unfortunately, this directional information turns out to be closely associated with the so-called *Real- μ problem*, i.e. the problem of calculating structured singular values for real -- rather than complex-valued -- plant modeling errors; this problem has been studied by Doyle and is generically very difficult. Its solution appears to be beyond the state of the art, at least in the near future.

The only reasonable alternative appears to be to translate the prior knowledge of structured uncertainty into an equivalent unstructured uncertainty. It is still a very hard problem to design a compensator with guaranteed performance characteristics in the presence of these modeling errors. We have transformed the problem into what Doyle calls the *μ -synthesis problem*, which unfortunately is also very hard to solve.

A promising theoretical and algorithmic approach to the solution of the *μ -synthesis*

problem has been developed. The theory utilizes the use of Hankel norms in approximating L_∞ functions using H_∞ functions. Certain procedures have been developed which would indicate whether or not the posed performance specifications are "too tight" for the level of modeling error present. In this case, the control system designer will have to relax the performance specifications, typically expressed as bounds on the sensitivity function maximum singular value, over some frequency ranges.

Maintaining stability in the presence of uncertainty has long been recognized as a crucial requirement for the closed-loop system. Classical designers developed the concepts of phase and gain margin to describe stability-robustness. In the modern control era, conditions for maintaining closed-loop stability in the presence of a single, unstructured (i.e. norm-bounded) modeling uncertainty have been formulated in terms of a singular value inequality on the closed-loop transfer function. It is only recently that the issue of multiple modeling uncertainties appearing at different locations in the feedback loop and the related requirement of performance-robustness have been addressed. Multiple unstructured uncertainty blocks, parameter uncertainty, and performance specifications give rise to so-called structured uncertainty. A new analysis framework based on the structured singular value has been developed by J. Doyle to assess the stability and performance robustness of a linear, time-invariant (LTI) feedback system in the presence of structured uncertainty. The structured singular value μ yields a necessary and sufficient condition for robust stability and performance.

While the analysis aspect of LTI feedback design is well-established, the μ -synthesis problem remains open. The purpose of this research has been to develop a practical methodology (based on μ) for the synthesis of robust feedback systems. That is, the design process will ensure the resulting feedback system is stable and performs satisfactorily in the event the actual physical plant differs from the design model (as it surely will). The motivation for an alternative to D,K iteration is due to the nonconvex nature of the μ -synthesis problem. Nonconvexity may lead to local minima, therefore it is essential that several independent methods be available to examine the problem.

This research has produced a new approach to the design of LTI feedback systems. For a given plant, the Youla parameterization describes all stabilizing compensators in terms of a stable, causal operator Q . LTI feedback design may be viewed as simply a procedure for choosing the appropriate Q to meet certain performance specifications. Thus, the design process imposes two constraints on the free parameter Q : (1) stability and causality (i.e. Q must be an H_∞ function); (2) Q must

produce a closed-loop system that satisfies some performance specification. The design objective of interest here is performance robustness, which can be stated in terms of a frequency domain inequality using the structured singular value.

The CRM initially lifts the restriction of compensator causality and the synthesis problem with uncertainty is examined at each frequency. A feasible set of Q 's in the space of complex matrices satisfying the performance specification is constructed. Causality is then recovered via an optimization problem which minimizes the Hankel norm (i.e. the measure of noncausality) of Q over the feasible set. If the problem is well posed (i.e. the performance specifications are not too stringent given the amount of modeling uncertainty), the resulting compensator nominally stabilizes the feedback system and guarantees robust stability and performance.

The theoretical foundation for the methodology have been established. Next, a research algorithm was written so that we can obtain numerical results. It was applied to two design examples to demonstrate its effectiveness. Excellent robust performance was obtained. However, the current generation of our CRM algorithms require very extensive *off-line* computational resources, because of the several optimization problems that must be solved to design the robust compensator.

1.6 Adaptive Redesign Strategies Following Failures

It is important to develop both high level (symbolic) and low level(quantitative) strategies for coping with control surface failures in submarines and aircraft. To compensate for a control surface failure, sufficient redundancy in the control authority must be provided by other control surfaces, thrust and moment producing mechanisms. To understand these issues, presently configured aircraft provide an opportunity for the development of such strategies.

Control failures in aircraft are not uncommon. Military aircraft can expect frequent damage to their control surfaces from enemy fire. However, even civil aircraft undergo such failures. A brief survey in [21] yielded almost 30 cases in which there were failures of controls other than engines. In all but five of these incidents, such malfunctions resulted in crashes, and loss of life to passengers and crew. In about half of these cases, the flight could have ended safely if the pilot had acted in a correct and timely manner; unfortunately, present procedures and training are inadequate to prevent many such accidents because corrective action must be taken extremely fast. What is needed is an automated means of helping the pilot to utilize the implicit multivariable redundancy of his many surfaces and thrust producing mechanisms so as to recover positive control of the aircraft.

The recently completed Ph.D. thesis of E. Wagner [21], under the supervision of

Professor Valavani, has made important strides toward the development of an on-board automated aid advisory for a C-130 aircraft. A rule-based expert system was developed to handle elevator-jam failures for the C-130 aircraft and its value illustrated using extensive simulations. This expert system produces an intelligent guide to pre-simulations of alternative controls (elevator tab, collective ailerons, symmetric flaps and engine thrust) using a high fidelity model of the aircraft. Pre-simulation of a recovery strategy was crucial because (a) often even a few degrees of available deflections could make all the difference, and (b) side-effects of doing the wrong thing could be devastating. The rule-based system was programmed using the OPS5 program.

This line of research is continuing on the on-going doctoral thesis of D. Obradovic [22] under the supervision of Professor Valavani. This research seeks to the development of alternative theoretical approaches to the control redesign problem which can be blended in a high-level symbolic system as described in [21].

2. PROGRESS IN SUBMARINE CONTROL SYSTEM DESIGNS

As stated in the original proposal, we were interested in the multivariable control of submarines so as to make a judgment on whether or not advanced adaptive control techniques are necessary for high-performance submarine control systems. Three Engineer's theses were completed during this period dealing with the submarine control problem. All three theses were written by Navy officers studying at MIT; one of them actually has served as a human controller in attack submarines. The theses investigated the design of multivariable control systems for submarines whose dynamics approximate those in the 688 class of attack submarines. We used dynamic coordination of the rudder, bow plane, and a split stern plane so as to provide independent roll control. The control system was designed so as to follow independent commands in desired pitch angle, inertial depth-rate, yaw-rate, and roll angle. This provides the potential for precise control of severe and demanding maneuvers, especially at high speeds. See publications [6], [12], [13], and [15].

Our research has demonstrated that active roll control has very beneficial effects. Its wise use allows the submarine to make tight high speed turns (in excess of those currently allowed under human control, due to safety considerations) with small depth excursions. In the absence of active roll control, the submarine can lose significant depth during these maneuvers.

Our research has also demonstrated that our designs were quite robust to the changing submarine dynamics as the speed varied. At most one needs to integrate the LQG/LTR designs (see publications [5], [16]) with a simple adaptive gain-scheduling algorithm, where the compensator gains are changed as a function

of speed. We do not believe that it is necessary to use more advanced adaptive methods for the multivariable control of submarines. Actually, a significant byproduct of our research relates to the saturation of the bow-plane in severe maneuvers. In this case, in order to maintain performance, it becomes necessary to design a restructurable control system. We did this by adapting the so-called "anti-reset windup" methodology to the multivariable case. More work along these lines is necessary before we have an integrated design methodology for adaptive restructurable control systems.

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LABORATORY FOR INFORMATION AND DECISION SYSTEMS
RESEARCH ON ROBUST STOCHASTIC ADAPTIVE CONTROL
ONR CONTRACT N00014-82-K-0582 (NR 606-003)

List of Publications

September 1, 1982 to December 31, 1987

- [1]. C.E. Rohrs, L.Valavani, M. Athans, and G. Stein, "Robustness of adaptive control algorithms in the presence of unmodeled dynamics," *Proc. 21st IEEE Conf. on Decision and Control*, Orlando, Florida, December 1982, pp. 3-11.
- [2]. L.Valavani and M. Athans, " Some critical questions about deterministic and stochastic adaptive control algorithms," *Proc. MELENCON 83*, Athens, Greece, May 1983.
- [3]. J. Krause, "Adaptive Control in the presence of high-frequency modeling errors," SM Thesis, LIDS-TH-1314, MIT, May 1983.
- [4]. J. Krause, M. Athans, S.S. Sastry, and L. Valavani, " Robustness studies in adaptive control," *Proc. 22nd IEEE Conf. on Decision and Control*, San Antonio, Texas, December 1983, pp. 977-981.
- [5]. M. Athans, " The relationship between the Zames representation and LQG compensators," *Proc. 22nd IEEE Conf. on Decision and Control*, San Antonio, Texas, December 1983, pp. 631-633.
- [6]. K.A. Lively, " Multivariable control system design for a submarine," SM Thesis, LIDS-TH-1379, MIT, May 1984.
- [7]. D. Orlicki, L. Valavani, M. Athans, and G. Stein, " Adaptive control with variable dead-zone nonlinearities," *Proc. American Control Conference*, San Diego, California, June 1984, pp. 1893-1898.
- [8]. C.E. Rohrs, M. Athans, L. Valavani, and G. Stein, " Some design guidelines for discrete-time adaptive controllers," *Proc. IFAC World Congress*, Budapest, Hungary, July 1984.
- [9]. L. Valavani and M. Athans, " On robust adaptive control algorithms," *Proc. DIGITECH*, Patras, Greece, July 1984.

- [10]. C.E. Rohrs, L. Valavani, M. Athans, and G. Stein, " Some design guidance for discrete-time adaptive controllers," *AUTOMATICA*, Vol. 20, 1984, pp. 653-660.
- [11]. D.M. Orlicki, " Model-reference adaptive control systems using a dead-zone nonlinearity," Ph.D. Thesis, MIT, April 1985.
- [12]. J.A. Mette, " Multivariable control of a submarine using the LQG/LTR method," Engineer thesis, MIT, May 1985.
- [13]. R.J. Martin, " Multivariable control for a submarine using active roll control," Engineer thesis, MIT, May 1985.
- [14]. C.E. Rohrs, L. Valavani, M. Athans, and G. Stein, " Robustness of continuous time adaptive control algorithms in the presence of unmodelled dynamics," *IEEE Trans. on Automatic Control*, Vol. AC-30, 1985, pp. 881-889.
- [15]. R.J. Martin, L. Valavani, and M. Athans, "Multivariable control of a submersible using the LQG/LTR design methodology", *Proc. American Control Conference*, Seattle, WA, June 1986; also accepted for publication to the IEEE Transactions on Automatic Control.
- [16]. M. Athans, " A tutorial on the LQG/LTR method", *Proc. American Control Conference*, Seattle, WA, June 1986.
- [17]. D. Milich, L. Valavani, and M. Athans, "Feedback system design with an uncertain plant", *Proc. 25th IEEE Conference on Decision and Control*, Athens, Greece, December 1986
- [18]. R.O. LaMaire, "Robust time and frequency domain estimation methods in adaptive control," Ph.D. Thesis, MIT, May 1987
- [19]. R.O. LaMaire, L. Valavani, M. Athans, and G. Stein, "A frequency domain estimator for use in adaptive control systems," *Proc. American Control Conference* Minneapolis, MN, June 1987
- [20]. D. Milich, "A methodology for the synthesis of robust feedback systems," Ph.D. Thesis, MIT, February 1988
- [21]. E. Wagner, "On-board automatic aid and advisory for pilots of control-impaired aircraft", Ph.D. Thesis, MIT, February 1988

[22]. D. Obradovic, "Adaptive methods for control system redesign," Ph.D. Thesis (in progress)