


IEEE Control Systems Magazine welcomes suggestions for books to be reviewed in this column. Please contact either Michael Polis or Zongli Lin, associate editors for book reviews.



\mathcal{L}_1 Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation

BY NAIRA HOVAKIMYAN and CHENGYU CAO

Reviewed by FRANK L. LEWIS

SIAM, 2010, ISBN: 978-0-898717-04-4, 330 pages, US\$85.

Controllers for linear systems can be analyzed for their performance and robustness in the frequency domain by using Bode plots, the Nyquist plot, the Nichols chart, and other tools. In the frequency domain, transient response can be analyzed in terms of bandwidth and resonant frequency properties, steady-state behavior is analyzed in terms of low-frequency gain, and robustness is analyzed in terms of gain, phase, and time-delay margins. Feedback control transfer functions can be selected to trade off performance and robustness. Frequency-domain analysis is highly developed in many application areas, including aircraft control [1], where there are strict performance requirements in terms of speed of response and overshoot, robustness requirements in the presence of varying or unknown parameters and disturbances, and requirements to limit control bandwidth due to the presence of vibratory system modes.

In aircraft control, the reliability of linear systems design has led to gain scheduling for flight controllers, which allows the design of controllers with guaranteed performance and robustness in prescribed regions of the flight envelope. Gain scheduling requires linear models, which are obtained using extensive and expensive wind tunnel testing.

Research in adaptive flight control started in the 1950s with a thrust to develop autopilots for supersonic aircraft with large flight envelopes while reducing the amount of gain-scheduling of linear controllers. Adaptive controllers were developed to cope with unknown dynamics, time-varying parameters, and unknown disturbances. There was intense early interest but also several failures of adaptive controllers [2]. Early developments in adaptive control were not sufficiently supported by theoretical analysis. Subsequent applications of Lyapunov theory and other techniques allowed rigorous analysis of adaptive controllers, resulting in modifications in parameter tuning laws, enhancements such as projection methods and robustifying terms, and guarantees on the ultimate boundedness of estimation errors and tracking errors. Tuning of adaptive controllers remains difficult, with parameters often being tuned by trial-and-error or simulation. Transient response and control signals do not have guaranteed bounds, and it is difficult to obtain robustness guarantees. The result is a continued reliance on gain-scheduled controllers.

In the 1980s the aerospace industry was transformed by the advent of fly-by-wire digital avionics systems. These systems improved the ability to control aircraft using digital controllers. Yet new configurations of aircraft are aerodynamically unstable, have more significant nonlinearities, and are required to have larger flight envelopes. Particularly difficult are the supermaneuverability requirements, which demand flight at high angles of attack, where linear approximation models are not valid. Military and aging commercial aircraft applications require the design of control systems that can recover from structural or actuator failures. These developments generated a renewed interest in adaptive control, with adaptive controllers often used to augment gain-scheduled controllers. This combination affords some possibility of combining proven linear system design methods with adaptive control solutions.

Model reference adaptive controllers (MRAC) provide an approach to combining linear control design methods with adaptive controllers, since linear design techniques can be used to design the reference models so that they conform to MILSPEC requirements [3] and provide performance improvement when tracking errors are small. However, bounds on the transient behavior of control signals cannot be given, and the designed controls are often

high gain and can contain high frequency modes that excite system vibratory modes. In MRAC, tuning adaptation parameters remains a challenge, and control signal oscillations can occur with improperly tuned gains. This behavior is related to the certainty-equivalence nature of these controllers, which use the estimated parameters as control gains. From the estimation perspective, faster estimation is needed, but larger adaptation gains make the system more sensitive to time delays since, as the adaptation gains are increased to reduce errors, the phase margin tends to zero. That is, attempts to improve performance result in decreased robustness. This tradeoff occurs because the compromise between performance and robustness is relegated to the tuning of the adaptation gains involved in nonlinear adaptation laws.

\mathcal{L}_1 ADAPTIVE CONTROL

The book *\mathcal{L}_1 Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*, by Naira Hovakimyan and Chengyu Cao, describes a new method known as \mathcal{L}_1 adaptive control for bringing linear system design methods into adaptive control to confront several open problems, including the selection and tuning of adaptation gains, guaranteed transient performance of both states and controls, and guaranteed robustness margins. Fast adaptation to allow aggressive maneuvering and recovery from failures is also confronted. In standard MRAC, the control signals are given in terms of parameter estimates generated by adaptive tuning algorithms. The control is of the form $u(t) = \hat{\eta}(t)$, where the signal $\hat{\eta}(t)$ depends on the parameter estimates and reference trajectory. On the other hand, in \mathcal{L}_1 adaptive control, the control is of the form $u(s) = D(s)\hat{\eta}(s)$, where $D(s)$ is a design filter. A second key feature in \mathcal{L}_1 adaptive control is the use of a state-prediction system, whose dynamics can be prescribed and provide the reference model to be followed when the tracking and estimation errors are small. The third feature is the tuning of parameter estimates using projection methods. The block diagram of the resulting closed-loop adaptive system consists of linear components that can be analyzed using a mixture of linear frequency response techniques and adaptive control techniques.

The introduction of the design filter $D(s)$ up front results in a method that, on one hand, decouples adaptation from robustness and, on the other hand, tightly couples proven linear design methods with adaptive control, thus allowing for rigorous analysis of the properties of the closed-loop system. It is shown that the introduction of $D(s)$ allows bounds to be derived not only on the steady-state errors, but also on the L_∞ norms of the transient control signals and states. A transfer function that depends on the design filter $D(s)$ and the reference model is the key item in the analysis. The bounds are uniform with respect to time and depend on the reciprocal of the square root of the adaptation gain and the \mathcal{L}_1 norms of $D(s)$ and the largest values

Throughout the book there are simulation examples that confront real-world issues.

of the unknown parameters. For this reason, the method is named “ \mathcal{L}_1 adaptive control.” It is shown that, as the adaptation gains tend to infinity, the L_∞ norms of the controls, states, and estimation errors are bounded for all time. Moreover, the states and controls approach the states and controls of an ideal reference model that can be designed by selecting $D(s)$ and the state predictor dynamics. This feature allows loop shaping by selecting the predictor dynamics independently of the adaptation gains.

The bounds on the L_∞ norms of the states, controls, and error signals become smaller as the adaptation gains are increased. That is, the selection of very large adaptation gains, in the presence of the design filter $D(s)$, allows for fast adaptation and also small error bounds. This result, which is in contrast to MRAC, means that the \mathcal{L}_1 adaptive controller is easy to tune by appropriate selection of the filter, while setting the adaptation gains large. The design filter $D(s)$ limits the bandwidth of the control signals and hence can be selected to guarantee that the control signals are within the allowed bandwidth of the systems and do not excite high-frequency modes. Finally, as the adaptation gains are increased, it is proven that the closed-loop \mathcal{L}_1 adaptive controller has gain, phase, and time-delay margins that approach values that can be designed by the choice of the filter $D(s)$.

In summary, \mathcal{L}_1 adaptive control allows loop shaping and the design of robustness margins by selecting design filters, as well as good performance in terms of fast adaptation and convergence by selecting the adaptation gains large. This combination is accomplished by a controller design that incorporates linear filters up front and thus allows rigorous analysis using a combination of linear systems theory and nonlinear systems analysis techniques.

LAYOUT OF THE BOOK

Chapter 1 provides a historical overview and introduces \mathcal{L}_1 adaptive control by comparing it to MRAC using a simple example. Chapter 2 covers \mathcal{L}_1 adaptive control using state feedback for the case of matched uncertainties. It is emphasized that loop shaping can be accomplished by selecting design filters, while at the same time fast adaptation occurs by selecting large adaptation gains. Systems with unknown input gains are treated. It is shown that the limiting time-delay margin as the adaptation gain becomes large depends on the design filter and not on the adaptation gain. The gain margin is analyzed. The extension is made

to unmodeled actuator dynamics, and a simulation shows good performance on the Rohrs counterexample. The extension is further made to nonlinear systems. A thorough discussion is made of linear filter design for performance and robustness tradeoffs.

In Chapter 3 the \mathcal{L}_1 adaptive controller is designed for state feedback in systems with unmatched uncertainties and unknown nonlinearities and input gains. Nonlinear strict feedback systems are treated. The extension is made to multi-input, multi-output systems. An alternative estimation scheme is given that does not enforce matching conditions. Chapter 4 extends the design to output feedback. An output predictor is used along with a design filter. An \mathcal{L}_1 condition must hold on a transfer function that depends on the design filter. It is shown how to design to satisfy this condition for minimum-phase systems with relative degree 1 or 2. An \mathcal{L}_1 adaptive controller is designed for non-strictly-positive-real reference systems. Chapter 5 develops \mathcal{L}_1 adaptive controllers for time-varying reference systems. First, linear time-varying systems are considered and then nonlinear systems with unmodeled dynamics.

Throughout the book there are simulation examples that confront real-world issues, such as the Rohrs example, wing rock in aircraft, highly nonlinear systems, the two-cart benchmark, and systems with time-varying dynamics. Chapter 6 is devoted to applications in actual flight validations in systems relating to several military organizations. A flight validation was done at the Naval Postgraduate School on a testbed of unmanned aerial vehicle with a commercial autopilot. The commercial autopilot was augmented with the \mathcal{L}_1 adaptive control architecture. Controls during aggressive path following and control surface failures were validated. The implementation results reveal that \mathcal{L}_1 adaptive control enhances the performance of the commercial autopilot. For the control surface failure, the commercial autopilot goes unstable while the \mathcal{L}_1 adaptation maintains stability. A key feature of the \mathcal{L}_1 controller is that it can adapt quickly without sacrificing robustness. The Rohrs example was extended to the flight test environment and the results show that the \mathcal{L}_1 controller outperforms MRAC.

A validation was carried out with NASA on the AirSTAR flight test vehicle. The same control parameters were used across the flight envelope with no redesign and no gain scheduling. Only the predictor parameters were gain scheduled to conform to MILSPEC requirements at multiple points in the operating regime. Validations were carried out for high-angle-of-attack captures and sudden engine failures. In each case the \mathcal{L}_1 controller adapted quickly enough to preserve stability and meet the desired performance, as predicted by theory.

The appendices cover systems theory, parameter adaptation using projection operators, and linear matrix inequalities.

DISCUSSION

The book \mathcal{L}_1 Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation, by Naira Hovakimyan and Chengyu Cao, contains mathematical analysis that is at times complex, yet it is always easy to read. The presentation is a mixture of didactic development of the control architectures, followed by proofs of the controller properties. As such, the important concepts are kept in the flow of the text, while more abstruse concepts are relegated to the proofs for exploration by the determined reader. Many auxiliary variables and signals are used, and it is sometimes difficult to keep track of the many symbols, but at the appropriate point when the main design equations are presented, the reader is usually referred to where each symbol is defined. The performance bounds are often complex expressions, but trends are explained so that the reader can intuitively grasp the relationships of importance. The writing is clear and explanatory, and ideas are developed in a logical fashion. The simulation examples show how to apply the techniques to design \mathcal{L}_1 adaptive controllers, illustrating which of the parameters must be selected and which are simply assumed to exist, as in the case, for instance, of bounds on the maximum values of unknown parameters.

The development of \mathcal{L}_1 adaptive control was initially encouraged by discussions with the defense industry. It has been shown in numerous actual flight validations that, in the presence of large adaptation gains, \mathcal{L}_1 adaptive control can provide a decoupling between robustness and performance tradeoffs by selecting linear filters. The performance of an \mathcal{L}_1 -adaptive closed-loop system is predictable, and as such, it seems to be a viable alternative to gain-scheduled MRAC-augmented baseline controllers that are popular today in the aerospace industry.

REVIEWER INFORMATION

Frank L. Lewis spent six years in the U.S. Navy during 1971–1977, serving as navigator aboard the frigate USS *Trippe* (FF-1075) and executive officer and acting commanding officer aboard USS *Salinan* (ATF-161). He obtained the Ph.D. in 1981 at the Georgia Institute of Technology in Atlanta, where he was employed as a professor from 1981 to 1990. He is a professor of electrical engineering at the University of Texas at Arlington, where he was awarded the Moncrief-O'Donnell Endowed Chair in 1990 at the Automation and Robotics Research Institute.

REFERENCES

- [1] B. L. Stevens and F. L. Lewis, *Aircraft Control and Simulation*, 2nd ed. New York: Wiley, 2003.
- [2] B. Anderson, "Failures of adaptive control theory and their resolution," *Commun. Inform. Syst.*, vol. 5, no. 1, pp. 1–20, 2005.
- [3] "Flying qualities of piloted aircraft," U.S. Dept. Defense Military Specification, MIL-STD-1797B, Feb. 2006.