Preface

This book gives a comprehensive overview of the recently developed $\mathcal{L}_1$ adaptive control theory with detailed proofs of the fundamental results. The key feature of $\mathcal{L}_1$ adaptive control architectures is the guaranteed robustness in the presence of fast adaptation. This is possible to achieve by appropriate formulation of the control objective with the understanding that the uncertainties in any feedback loop can be compensated for *only* within the bandwidth of the control channel. By explicitly building the robustness specification into the problem formulation, it is possible to decouple adaptation from robustness via continuous feedback and to increase the speed of adaptation, subject only to hardware limitations. With $\mathcal{L}_1$ adaptive control architectures, fast adaptation appears to be beneficial both for performance and robustness, while the trade-off between the two is resolved via the selection of the underlying filtering structure. The latter can be addressed via conventional methods from classical and robust control. Moreover, the performance bounds of $\mathcal{L}_1$ adaptive control architectures can be analyzed to determine the extent of the modeling of the system that is required for the given set of hardware.

The book is organized into six chapters and has an appendix that summarizes the main mathematical results, used to develop the proofs.

Chapter 1 starts with a brief historical overview of adaptive control theory. It proceeds with an introduction to the main ideas of the $\mathcal{L}_1$ adaptive controller. Two equivalent architectures of model reference adaptive controllers (MRAC) are considered next, and the challenges of tuning these schemes are discussed. The chapter proceeds with analysis of a stable scalar system with constant disturbance and introduces the main idea of the $\mathcal{L}_1$ adaptive controller. Two key features are analyzed in detail: the closed-loop system's guaranteed phase margin and the uniform bound for its control signal.

Chapter 2 presents the $\mathcal{L}_1$ adaptive controller for systems in the presence of matched uncertainties. It starts from linear systems with constant unknown parameters and develops the proofs of stability and the performance bounds. The results in this section prove that the $\mathcal{L}_1$ adaptive controller leads to guaranteed, uniform, and decoupled performance bounds for both the system's input and output signals. First, it is proved that with fast adaptation the state and the control signal of the closed-loop nonlinear $\mathcal{L}_1$ adaptive system follow the same signals of a reference linear time-invariant (LTI) system for all $t \geq 0$. As compared to the original reference system of MRAC, which assumes perfect cancelation of uncertainties, the reference LTI system in $\mathcal{L}_1$ adaptive control theory assumes only partial cancelation of uncertainties, namely, those that are within the bandwidth of the control channel. Despite this, and because it is an LTI system, its response scales uniformly with changes in the initial conditions, reference inputs, and uncertain parameters. Therefore, the response of the closed-loop nonlinear $\mathcal{L}_1$ adaptive system also scales with all the changes in initial conditions, reference inputs, and uncertainties. Next, it is proved that this LTI reference system can be designed to achieve the desired control specifications. This step is the key to the trade-off between performance and robustness and is reduced to tuning the structure and the bandwidth of a stable strictly proper bandwidth-limited linear filter. Thus, the complete performance bounds of the nonlinear $\mathcal{L}_1$ adaptive controller are presented via two terms: the first is inversely proportional to the rate of adaptation, while the second depends upon the bandwidth of a linear filter. This decoupling between adaptation and robustness is the key feature of the $\mathcal{L}_1$ adaptive controller. The chapter proceeds by extending the class of systems to accommodate an uncertain system input gain, time-varying parameters, and disturbances. A rigorous proof for a lower bound of the time-delay margin of the closed-loop $\mathcal{L}_1$ adaptive system is provided in the case of open-loop linear systems with unknown constant parameters. This lower bound is computed from an LTI system, using its phase margin and its gain crossover frequency. The loop transfer function of this LTI system has a decoupled structure, which allows for tuning its phase margin or time-delay margin via the selection of the bandwidth-limited filter. The chapter proceeds by considering unmodeled actuator dynamics and nonlinear systems in the presence of unmodeled dynamics and uses the well-known Rohrs' example to provide further insights into the $\mathcal{L}_1$ adaptive controller. Other benchmark applications are discussed. An overview of tuning methods for the design of this filter for a performance–robustness trade-off is presented toward the end, and as an example, an LMI-based solution is described with certain (conservative) guarantees.

Chapter 3 extends the $\mathcal{L}_1$ adaptive controller to accommodate unmatched uncertainties. It starts with nonlinear strict-feedback systems, for which the $\mathcal{L}_1$ adaptive backstepping scheme is presented. The chapter proceeds with an extension to multi-input multi-output (MIMO) nonlinear systems in the presence of general unmatched uncertainties and unmodeled dynamics or, alternately, unknown time- and state-dependent nonlinear cross-coupling, which cannot be controlled by recursive (backstepping-type) design methods. Two different adaptive laws are introduced, one of which, being piecewise constant, is directly related to the sampling parameter of the CPU. There are certain advantages to this new type of adaptive law. It updates the parametric estimate based on the hardware (CPU) provided specification. At the sampling times, the adaptive law reduces one of the components of the identification error to zero, with the residual being proportional to the sampling interval of integration. This implies that by increasing the rate of sampling, one can reduce the influence of the residual term on the performance bounds. The uniform performance bounds are derived for the control signal and the system state as compared to the corresponding signals of a bounded closed-loop reference system, which assumes partial cancelation of uncertainties within the bandwidth of the control channel. This MIMO architecture has been applied to NASA's Generic Transport Model (GTM), which is part of the AirSTAR system, and to Boeing's X-48B blended wing body aircraft. Appropriate references are provided.

Chapter 4 presents the output feedback solution. It starts by considering first-order reference systems for performance specifications. Next, it considers more general reference systems that do not verify the SPR property for their input-output transfer function. In the second case, the piecewise-constant adaptive law is invoked for compensation of the effect of uncertainties on the system's regulated output. Similar to state-feedback architectures, a closed-loop reference system is considered, which assumes partial cancelation of uncertainties within the bandwidth of the control channel. However, unlike the state-feedback case, the sufficient condition for stability in this case couples the system uncertainty with the desired reference-system behavior and the filter design. The two-cart benchmark example is discussed as an illustration of this extension. Also, the flight tests at the Naval Postgraduate School are based on the solutions from this chapter.

Chapter 5 presents an extension to accommodate linear time-varying (LTV) reference systems. This extension is critical for practical applications. For example, in flight control, quite often the performance specifications across the flight envelope are different at different operational conditions. This leads to a time-varying reference system, the analysis of which cannot be captured by the tools developed in previous chapters. Appropriate mathematical tools for addressing this class of systems are presented in the appendices. The chapter also presents a complete solution for nonlinear systems in the presence of unmodeled dynamics. The uniform performance bounds of the system state and the control signal are computed with respect to the corresponding signals of an LTV reference system, which meets different transient specifications at different points of the operational envelope.

Chapter 6 summarizes some of the further extensions not captured within this book, gives an overview of the applications and the flight tests that have used this theory, and states the open problems and the challenges for future thinking. Appropriate references are provided.

The book concludes with appendices, where basic mathematical facts are collected to support the main proofs.

The book can be used for teaching a graduate-level special-topics course in robust adaptive control.

Notations

The book interchangeably uses time-domain and frequency-domain language for representation of signals and systems. For example, $\xi(t)$ and $\xi(s)$ denote the function of time and its Laplace transform, respectively. However, this should not confuse the reader, as all the equations in the book are written using only one argument, either t or s. There are no mixed notations in any of the equations of the book. By smoothly moving from one form of representation to another, we streamlined the analysis and proofs. Whenever needed, thorough explanations are provided. Unless otherwise noted, $\|\cdot\|$ will be used for the 2-norm of a vector. Finally, $\mathfrak{L}(\xi(t))$ is used to denote the Laplace transform of the time signal $\xi(t)$.

Acknowledgments

The most valuable performance metric for every feedback solution is the experimental validation of it. We are especially grateful to Isaac Kaminer and Vladimir Dobrokhodov from the Naval Postgraduate School for sharing their challenging problems with us and for their insights into the performance of the solutions. Their numerous successful flight tests in the context of very challenging multiconstraint applications in Camp Roberts, CA, since 2006 have built strong credibility in this theory. The invaluable insights shared with us by Kevin Wise and Eugene Lavretsky from Boeing and by Brett Ridgely and Richard Hindman from Raytheon lie behind the assumptions, solutions, and techniques in this book. Kevin Wise and Eugene Lavretsky motivated many of the problems and posed a number of *good* *questions* to us. Kevin provided tremendous encouragement, and if not for his efforts of transitioning adaptive control into industry products and his feedback from that experience, this theory would not have been developed. We are grateful to Randy Beard from Brigham Young University, Jonathan How from MIT, and Ralph Smith from North Carolina State University for their experimental results in different applications.

We would like to acknowledge a number of research scientists, postdoctoral fellows, and graduate students who have worked with us at different times and whose Ph.D. dissertations and technical papers contribute to the chapters in this book. Among these, we are especially thankful to Vijay Patel of the Indian Ministry of Defense for his help and support with flight control applications during his sabbatical stay at Virginia Tech in the initial stages of the theory's development. Graduate students Jiang Wang, Dapeng Li, Enric Xargay, Zhiyuan Li, Evgeny Kharisov, Hui Sun, and Kwang-Ki Kim contributed to the development of the results in this book while working on their dissertations. We are profoundly grateful to Evgeny Kharisov for his support with the editing of the book and his careful check of all the proofs and examples and for generating a large number of new simulations to support the theoretical claims. We would like to acknowledge Tyler Leman for his efforts in transitioning the theory to the X-48B platform and Alan Gostin for developing the MATLAB-based toolbox in support of this theory. Enric Xargay's contributions to the flight tests of NASA's GTM (AirSTAR) model deserve special mention. His feedback from that experience significantly enhanced the understanding of the theory. Enric's help with the last proofreading of the book was extremely useful and constructive.

We had invaluable comments from Karl J. Åström regarding the structure and the layout of the book. His recommendations regarding the additional simulations and examples helped to provide deeper insights for the reader. He also helped us with the editing of the historical overview. We are grateful to Roberto Tempo for numerous useful discussions on the design and the optimization problems of the methods in this book. Geir Dullerud from the University of Illinois at Urbana-Champaign provided a number of useful comments. Several discussions with Allen Tannenbaum and Tryphon Georgiou have influenced our thinking over the years. We are thankful to Xiaofeng Wang for extending the theory to event-triggered networked systems and to Tamer Başar for his interest in quantization of the $\mathcal{L}_1$ adaptive controller.

Irene Gregory from the NASA Langley Research Center (LaRC), Dave Doman and Michael Bolender from the Wright-Patterson Air Force Research Laboratory (AFRL), and John Burken from the NASA Dryden Flight Research Center have shared different mid- to high-fidelity models with us, explaining the particular challenges and objectives of those applications. We would like to thank Irene with a special mention for her efforts in transitioning the theory to the flight tests on the GTM at NASA LaRC. We also would like to acknowledge Kevin Cunningham, Austin Murch, and the rest of the staff of the AirSTAR flight test facility for sharing their insights into flight dynamics and for their support with implementation of the control law on the GTM. Our interactions with Johnny Evers from the Eglin AFRL, Siva Banda from the Wright-Patterson AFRL, and the research groups they lead were always full of inspiration. The financial support for this research was initially provided by the Air Force Office of Scientific Research, and it was leveraged into various programs across the country, including *Certification of Advanced Flight Critical Systems* at Wright-Patterson AFRL and NASA's *Integrated Resilient Aircraft Control*.

Last, but not least, we would like to thank our families for their unconditional dedication, love, and support, and to whom—with our humble gratitude—we dedicate this book.

Champaign, Illinois
Storrs, Connecticut

Naira Hovakimyan
Chengyu Cao

March 2010