Preface

Both *control of PDEs* (partial differential equations) and *adaptive control* are challenging fields. In control of PDEs the challenge comes from the system dynamics being infinite-dimensional. In adaptive control, the challenge comes from the need to design feedback for a plant whose parameters may be highly uncertain (much more than a few percent or a few dozen percent relative to a nominal value) and which may be open-loop unstable, requiring simultaneous control and learning. This requirement, in turn, results in adaptive controllers being nonlinear even when the plants are linear. Because of this, the detailed behavior of adaptive feedback systems (for example, transient performance) is very hard to predict.

With challenges come controversies. The 1980s were a decade of controversies about adaptive control, where a community, which was used to the predictability of linear feedback systems and to the availability of linear performance limitation results like Bode's, was faced perhaps for the first time with a feedback problem that was fundamentally nonlinear. The controversies were fueled not only by the lack of understanding of particular adaptive feedback schemes, but by the lack of appreciation for the nature of the problem itself—stabilization in the presence of large parametric uncertainty. The controversies somewhat subsided in the early 1990s when new adaptive schemes, based on the method of "integrator backstepping," emerged with which one can systematically reduce the performance bounds (in L_∞ and L_2 norms). However, the unpredictability of adaptive controllers (in a detailed sense, rather than in a "bulk," performance bound sense) in the absence of "persistency of excitation," namely, the dependence of the transient and asymptotic behaviors of adaptive feedback schemes on initial conditions, is fundamental to the problem. One could say that this difficulty is the point where the engineering appeal of adaptive control starts coming into question and its mathematical beauty begins. Despite its stormy adolescent period, adaptive control continues to be strong in its middle-aged years, with plenty of intellectual mileage left, and with a growing number of industrial applications, including even the fields where it had traditionally been considered to be "too risky," such as aerospace systems and flight control.

In the case of control of PDEs, no controversy of the same magnitude has existed, however, this field has chronically suffered from the inability to enthuse engineers.

The first reason is, again, the mathematical difficulty associated with PDEs. The second is the belief that finite-dimensional control design tools should suffice, since many (or most, though not all) PDEs are dominated by some finite-dimensional dynamic behavior. Hence, some form of model reduction, Galerkin approximation, or spatial discretization should suffice. If it only were this simple. If it were—the control engineers would be self-sufficient in the area of control of PDEs, starting with a quick step of reducing the problem to an ODE and then applying a finite-dimensional design. Unfortunately, neither is this step quick—in some cases it involves techniques that researchers in other fields spend entire careers honing, such as in fluid mechanics and solid mechanics—nor is the rest easy (and mathematically straightforward) even if the designer has successfully completed this first step of PDE→ODE reduction.

The question of approximating a PDE by an ODE for the purpose of control design is the same "chicken and egg" question as the question of model (dimension) reduction for control design for ODEs. One can never be sure that model reduction for the plant (based on open-loop considerations) will lead to a control design (performed on the reduced model) which will be effective for the original (unreduced) plant. For PDEs, this question translates into: how to perform spatial discretization in such a way that, upon finite-dimensional control design, and upon refinement of the discretization grid, the control gains converge as the spatial discretization step goes to zero. This is a highly nontrivial question. Very intuitive, "obvious" and standard discretization techniques used for PDE *simulation* (without control), lead to divergent control gain kernels upon grid refinement. It often takes completely non-obvious discretization ideas to achieve convergence of gains (though convergence of signals may be achieved even when the gains are not convergent). The moral of this story is that a priori approximation of a PDE by an ODE, for the purpose of feedback design, is not a problem that is solvable in general. One can easily approximate the PDE after the controller has been designed, but it is not safe to do so for the purpose of designing a controller.

Hence, it is crucial to develop control tools in the continuum domain, for the PDEs, and invariant of model reduction choices.

* * *

What Does the Book Cover? This book undertakes a major task of solving adaptive control problems for PDEs. The difficulty of this task is even greater than the sum of the difficulties of the two components, namely, (1) control of PDEs and (2) adaptive control (for ODEs).

Indeed, until a few years ago, most researchers would have considered adaptive control of infinite-dimensional systems (and PDEs in particular) almost a misnomer because adaptive control requires an explicit parametrization of the control gains (either directly, or indirectly, through the plant parameters). Within the framework of classical control methods for PDEs, such as linear quadratic optimal control, this approach would literally require that an operator Riccati equation (infinite-dimensional and nonlinear, though algebraic, independent of time) be solved at each time step.

With the method of "backstepping," which we recently introduced for boundary control of PDEs, the need for solving a new operator Riccati equation at each time step, for each new parameter estimate, is eliminated. With the backstepping approach the control gains are explicitly parametrized in terms of both the spatial variables and the plant parameters. One only needs to run online estimation of the plant parameters. The computation of the gains is a matter of plugging the plant parameters into an explicit formula, or at most of solving a linear problem of a very easy kind such as an integral equation in the spatial variable (derived from a linear hyperbolic PDE which characterizes the control gains) at each time step.

This book focuses on adaptive control of *parabolic* PDEs. This is the most forgiving class of PDEs, at least in terms of modeling errors, because only a finite number of eigenvalues of the plant can be in the right half plane or on the imaginary axis. Hyperbolic PDEs require additional tools as compared to those presented in this book and contain peculiarities related primarily to these equations being second-order in time, rather than being infinite-dimensional. A successful and thorough development for hyperbolic PDEs will likely require years of additional research. For this reason we restrict our attention to parabolic PDEs.

The efforts in adaptive control of PDEs begun in the late 1980s and have continued through the 1990s. The leading authors in these efforts were Balas, Bentsman, Demetriou, Duncan, Kobayashi, Orlov, Pasik-Duncan, and a few other authors.

The backstepping approach for PDEs, which has started crystalizing for non-adaptive problems around the year 2002, allows us to make a major leap forward methodologically, and to design adaptive controllers for previously intractable classes of problems. The emphasis in the book is fivefold:

- unstable PDEs,
- boundary control (and boundary sensing),
- infinite relative degree problems,
- simultaneous lack of state measurement and parameter knowledge (adaptive output-feedback),
- functional unknown parameters (infinite-dimensional parameter vector).

We of course proceed pedagogically, tackling these problems one at a time, and putting them all together in one of the final chapters (Chapter 13).

On the methodological level, the book contains an ambitious and exhaustive catalog of approaches to adaptive control. Each of these approaches arises from different parametrization of the system and different methods for parameter identifier design. Our book mirrors the categorization of methods introduced in the second author's 1995 classic monograph [67]. Three methods for identifier design are employed,

- Lyapunov-based,
- passivity/observer-based,
- swapping based.

In addition, two parametrizations are used,

- u-model (the original plant model),
- w-model (the error system model, after the application of the backstepping transformation).

This large collection of methods and parametric models creates a large number of combinations, each of which is a possible adaptive control scheme. We of course don't pursue all of them, but we do pursue a representative subset of combinations.

The book is structured as follows. It has two major parts: Part I introduces the underlying nonadaptive controller and observer designs (with a highlight on the problems where the controller and observer gains can be found explicitly) and Part II presents the adaptive designs. Roughly, the first third of the book is non-adaptive, whereas the remaining two thirds are adaptive.

* * *

Who is the Book For? The book should be of interest to any and every researcher who has worked on adaptive control of ODEs. It presents surprisingly accessible solutions to adaptive control problems for PDEs, which one might expect to be abstract and dry. Quite on the contrary, a researcher, student, or engineer in adaptive control will find a wealth of examples that are solved and presented in full analytical and numerical detail, graphically illustrated, and with intuition provided that helps understand the results.

Researchers in PDEs and, more broadly, mathematics (including nonlinear functional analysis), will find the book filled with problems of a new kind. The problems abound in mathematical challenges that do not arise in any other field except control theory, and, more specifically, parameter-adaptive control. Adaptive control of *linear* PDEs is a truly unique stepping stone toward the challenging world of nonlinear PDEs. In adaptive control of linear PDEs the overall system is nonlinear, due to the presence of the parameter identifier dynamics and its state feeding into the controller. However, this nonlinearity is sufficiently benign, and tractable globally, unlike those that one encounters in problems like 3D Navier-Stokes equations or nonlinear reaction-diffusion systems that exhibit finite-time blowup.

In addition to engineers who are control specialists, some chemical engineers and process dynamics researchers should find the book of interest, as it covers a broad class of parabolic PDEs, which includes the reaction-advection-diffusion systems, with spatially-varying coefficients, which arise in process dynamics.

Likewise, mechanical and aerospace engineers dealing with various mechanics problems of parabolic kind (thermal, fluid, combustion, and various other diffusion-dominated problems), even if control is not their focus, should find the book of interest, since the methods employed are not from the generic control catalog (LQR, pole placement) but are very structure-specific and PDE-specific. Physicists should find the book of interest for the same reason.

Even a researcher with no interest whatsoever in control systems in general can enjoy the book on two levels:

1. The methods for control developed in this book "morph" the system dynamics in an uncommonly intuitive way. For example, these methods are capable of removing *spatially distributed*, domain-wide reaction effects using inputs only from the boundary of the PDE domain. This is only an example of a much more general capability of the method to radically alter the nature of the PDE dynamics within the domain, acting only through boundary conditions.
2. A component of the adaptive control design, which is de-emphasized in the book (for the sake of focus on control—but still significant), is parameter estimation, namely *system identification for PDEs*. A researcher who cares about this topic alone would find the book useful. We tackle problems with both constant and spatially-varying parameters, and both problems where the full state of the PDE is measured and where only a boundary value is measured (adaptive observers).

Regarding required background to read this book, relatively little is needed beyond calculus. The book is self contained in terms of most of the nonlinear stability analysis tools, special functions background, or functional analysis techniques and inequalities employed.

Acknowledgments. We would like to thank Ole Morten Aamo, Saverio Bolognani, and Yuri Orlov for their contributions in Chapters 6, 12, and 13. We sincerely appreciate Andrea Serrani's and Gang Tao's generous input on the book. We would like to thank Joseph Bentsman, Michael Demetriou, Tyrone Duncan, Hartmut Logemann, Bozenna Pasik-Duncan, Stewart Townley, and John Wen for their inspiring early work on system identification and adaptive control of infinite-dimensional systems. We appreciate Brian Anderson's and Steve Morse's encouraging feedback when we were making our first steps on the topic of adaptive control of PDEs in late 2004. We thank Laurent Praly and Irena Lasiecka for their interest and technical input on our work. We are grateful to Petar Kokotovic for his encouragement and support. We dedicate this book to our daughters Katerina, Victoria, and Alexandra, and our wives Olga and Angela.

La Jolla, California,
March 2009

Andrey Smyshlyaev
Miroslav Krstic

Contents