Control Theory for Physicists

Control theory, an interdisciplinary concept dealing with the behavior of dynamical systems, is an important but often overlooked aspect of physics. This book is the first broad and complete treatment of the topic tailored for physicists, one that goes from the basics right through to the most recent advances. Simple examples develop a deep understanding and intuition for the systematic principles of control theory, beyond the recipes given in standard engineering-focused texts. Up-to-date coverage of control of networks and complex systems and a thorough discussion of the fundamental limits of control, including the limitations placed by causality, information theory, and thermodynamics, are included. In addition, this text explores important recent advances in stochastic thermodynamics on the thermodynamic costs of information processing and control. For all students of physics interested in control theory, this classroom-tested, comprehensive approach to the topic with online solutions and further materials delivers both fundamental principles and current developments.

John Bechhoefer is Professor of Physics at Simon Fraser University and Fellow of the American Physical Society. His research is at the intersection of thermodynamics, statistical physics, control theory, and information theory.

Control Theory for Physicists

JOHN BECHHOEFER
Simon Fraser University

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

Shaftesbury Road, Cambridge CB2 8EA, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

103 Penang Road, #05–06/07, Visioncrest Commercial, Singapore 238467

Cambridge University Press is part of Cambridge University Press & Assessment, a department of the University of Cambridge.

We share the University's mission to contribute to society through the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781107001183

DOI: 10.1017/9780511734809

First published 2021

A catalogue record for this publication is available from the British Library

ISBN 978-1-107-00118-3 Hardback

Additional resources for this publication at www.cambridge.org/bechhoefer

Contents

Preface *page* ix

Part I Core Material 1

1 Historical Introduction 3
1.1 Historical Overview 4
1.2 Lessons from History 12
1.3 Control and Information 21
1.4 Notes and References 21

2 Dynamical Systems 24
2.1 Introduction: The Pendulum as a Dynamical System 25
2.2 General Formulation 28
2.3 Frequency Domain 36
2.4 Time Domain 46
2.5 Stability 52
2.6 Bifurcations 58
2.7 Summary 60
2.8 Notes and References 61
Problems 62

3 Frequency-Domain Control 67
3.1 Basic Feedback Ideas 68
3.2 Two Case Studies 75
3.3 Integral, Derivative, and PID 78
3.4 Feedforward 83
3.5 Stability of Closed-Loop Systems 88
3.6 Delays and Nonminimum Phase 92
3.7 Designing the Control 100
3.8 MIMO Systems 110
3.9 Summary 121
3.10 Notes and References 121
Problems 122

4 Time-Domain Control 129
4.1 Controllability and Observability 130

4.2	Control Based on the State	142
4.3	Control Based (Indirectly) on the Output	148
4.4	Summary	156
4.5	Notes and References	157
	Problems	157
5	**Discrete-Time Systems**	161
5.1	Discretizing Signals	163
5.2	Tools for Discrete Dynamical Systems	175
5.3	Discretizing Dynamical Systems	181
5.4	Design of Digital Controllers	186
5.5	Summary	196
5.6	Notes and References	197
	Problems	198
6	**System Identification**	205
6.1	Physics or Phenomenology?	206
6.2	Measuring Dynamics	207
6.3	Model Building	222
6.4	Model Selection	227
6.5	Model Reduction	233
6.6	Summary	236
6.7	Notes and References	237
	Problems	238
	Part II Advanced Ideas	249
7	**Optimal Control**	251
7.1	One-Dimensional Example	252
7.2	Continuous Systems	254
7.3	Linear Quadratic Regulator	259
7.4	Dynamic Programming	263
7.5	Hard Constraints	269
7.6	Feedback	276
7.7	Numerical Methods	286
7.8	Summary	288
7.9	Notes and References	289
	Problems	291
8	**Stochastic Systems**	297
8.1	Kalman Filter	298
8.2	Linear Quadratic Gaussian Control	313
8.3	Bayesian Filtering	319
8.4	Nonlinear Filtering	326
8.5	Why State Estimation Can Be a Hard Problem	340

8.6 Stochastic Optimal Control 345
8.7 Smoothing and Prediction 348
8.8 Summary 352
8.9 Notes and References 352
Problems 355

9 Robust Control 363
9.1 Robust Feedforward 365
9.2 Robust Feedback 371
9.3 Risk 376
9.4 Worst-Case Methods: The $\mathcal{H}_\infty$ Min-Max Approach 386
9.5 Summary 395
9.6 Notes and References 398
Problems 399

10 Adaptive Control 404
10.1 Direct Methods 407
10.2 Indirect Methods 413
10.3 Adaptive Feedforward Control 424
10.4 Optimal Adaptive Control 430
10.5 Neural Networks 437
10.6 Summary 443
10.7 Notes and References 444
Problems 446

11 Nonlinear Control 454
11.1 Feedback Linearization 455
11.2 Lyapunov-Based Design 466
11.3 Collective Dynamics 469
11.4 Controlling Chaos 476
11.5 Summary 480
11.6 Notes and References 480
Problems 481

Part III Special Topics 491

12 Discrete-State Systems 493
12.1 Discrete-State Models 493
12.2 Inferring States and Models 500
12.3 Control 508
12.4 Summary 511
12.5 Notes and References 512
Problems 513

13 Quantum Control 516
13.1 Quantum Mechanics 517

13.2 Three Types of Quantum Control 520
13.3 Physical Example 523
13.4 How Different Is Quantum Control? 526
13.5 Summary 528
13.6 Notes and References 530
Problems 531

14 Networks and Complex Systems 533
14.1 Networks in a Nutshell 535
14.2 From Dynamics to Graphs to Networks 540
14.3 Structural Controllability 541
14.4 Minimum Inputs Problem 543
14.5 Control Effort 546
14.6 Complex Systems 551
14.7 Summary 554
14.8 Notes and References 555
Problems 557

15 Limits to Control 561
15.1 Causal Limits 561
15.2 Information-Theoretic Limits 574
15.3 Thermodynamic Limits 588
15.4 Summary 608
15.5 Notes and References 610
Problems 613

References 619
Index 641

Preface

This book extends a tutorial I wrote on control theory (Bechhoefer, 2005). In both the article and this book, my goal has been "to make the strange familiar, and the familiar strange."[1] The *strange* is control theory – feedback and feedforward, transfer functions and minimum phase, $\mathcal{H}_\infty$ metrics and Z-transforms, and many other ideas that are not usually part of the education of a physicist. The *familiar* includes notions such as causality, measurement, robustness, and entropy – concepts physicists think they know – that acquire new meanings in the light of control theory. I hope that this book accomplishes both tasks.

If You Are an Experimental Physicist

Control techniques such as feedback can dramatically increase the performance of your experimental apparatus, and most experimentalists, by necessity, pick up a certain amount of control theory informally. The results are often ok, but with a bit of savvy, they can usually be made much better. In this book, you will learn how to set the parameters and structure of a control system to maximize – by your own measure – the performance of your experiment. This is *optimal control*.

If you have learned a bit about proportional-integral-differential (PID) control and think that is enough, this book is for you: Although PID control is indeed a good generic solution – especially if you know how to use the D – knowing something about the dynamics of the system can lead to much better control. Generic solutions are fine for the family sedan, but a race car needs to be tuned.

An instructive example is the atomic force microscope (AFM). The original instruments were built by experimental physicists and were based on PI control (or slight variations). Starting about 2005, people in the industry realized that more advanced techniques could be valuable, and the current commercial AFM controllers now employ the full panoply of techniques discussed in this book, with thousandfold increases in scanning rates.[2]

[1] Like many sayings, this one has a complicated history. Its origin is attributed to Novalis (Baron Friedrich von Hardenberg, 1772–1801), poet, philosopher, and a founder of German Romanticism. In an unpublished fragment, he wrote, "To romanticize the world is to make us aware of the magic, mystery and wonder of the world; it is to educate the senses to see the ordinary as extraordinary, the familiar as strange, the mundane as sacred, the finite as infinite" (translated in Beiser, 1998). Subsequent reexpressions have boiled this down and added the converse, ending up as the pithier phrase I quote.

[2] Interestingly, the involvement of control specialists in the AFM industry came from the "top down." The National Science Foundation, a leading source of support for control theory in the United States, was concerned about an overly mathematical bent to control-theory research and began actively encouraging researchers to explore new applications, of which AFM was a prominent example.

If You Are a Theoretical Physicist

From theoretical physicists, a remark about control theory that I have heard is, "That's just engineering and linear systems" – said, if not with a sneer, then with an implication that there is nothing fundamental to interest a theoretical physicist. I hope to convince you that such sentiments give short shrift to one of the great ideas of the twentieth century, worth knowing on its own and for the deeper understanding of physical systems that it leads to.

What is different about control theory? In a word, *purpose*. Control always has a goal, and control theory shows how to achieve that goal. By contrast, physical theories describe nature, and historically physicists have gone to great lengths to remove the subjective from their theories. Thus, to study control theory as a physicist is a bit like licking the forbidden fruit. We know we are not supposed to do it, but it does seem interesting. I hope to assuage such feelings of guilt, as much in the natural world is designed, either explicitly by people, or implicitly through evolution and natural selection. And even if you do not share this way of looking at the world, learning about a great idea and its implications is worthy in itself, and the contrast can teach you much about "nonpurposeful" dynamics.

Control theory also deepens our appreciation of physical systems. There has been a steady expansion of physics and physical methods to the study of complex systems and networks. In many of these systems (such as the earth's climate, as discussed in Chapter 1), the notion of feedback is important, and this book will help give perspective on what those feedbacks can mean (Chapter 14). And, as we will discuss in Chapter 15, there are fundamental limits set on control from thermodynamics and causality (as well as information theory), and the effort to understand the interplay among these topics has generated much recent effort, including in my own research. Notions such as Maxwell's demon and Szilard's engine have evolved from thought experiments to real ones, and those achievements have then stimulated much ongoing theoretical work where control theory plays an important role.

If You Are a Control Theorist or Engineer

If you have a professional interest in control, you have many fine books to choose from, ranging from those focused on particular applications for the engineer to more theoretical books written for the control theorist and applied mathematician. What a physicist can offer is a concise introduction that takes a broader view and devotes more attention to the foundations of topics than is usual in textbook engineering discussions. I have tried to sort through a vast field, epitomized by the 3,500-page, three-volume series by Levine (2011b,c,a), and produce a broad, coherent view of what control can mean. Also, physicists focus on foundational issues, and this book includes topics rarely seen in introductory control-theory texts, such as the relations between stochastic and worst-case approaches to robust control (Chapter 9) and the limitations of control due to causality, information theory, and thermodynamics (Chapter 15).

Keeping this book to a reasonable size led to a rather ruthless competition, and much valuable material had to be cut. I apologize for all the pretty and important topics – Nyquist stability criterion! – that are not present.

Why This Book?

University libraries have shelves full of control-theory books. Why write another? Existing control-theory books are mostly written either for engineers or for control theoreticians and mathematicians. For a physicist, neither type of book is quite right. The mathematical approach is usually concise but not great at shaping intuitions. The engineering approach uses different notation (j for – i, Laplace transforms rather than Fourier transforms, factors of 2π, etc.) that can render familiar expressions mysterious. Moreover, the engineering style is quite different, with more emphasis on numerical, practical examples. Physicists prefer scaled variables and examples that are stylized to illustrate concepts. Here, I follow physics conventions by default but compromise to connect to control-theory texts.

This book has broader scope than most engineering texts. Because control theory is an essential part of an engineering degree, introductory texts are broad but low level, while advanced texts tend to be narrow. A similar problem exists with mathematics: Most physics departments have one or more courses in "mathematics for physicists" that compress half-a-dozen courses (tensors, boundary-value problems, partial differential equations, differential geometry, algebra, group theory, etc.) into a single course. You can view this book similarly (hence its title). Our topics are distilled from what is often covered in courses on classical, modern, digital, optimal, stochastic, robust, adaptive, and nonlinear control. Indeed, each chapter can be – and often is – the subject of a semester-long course. What physicist has time to take so many courses on control? Yet the answer is not the de facto current situation, where physicists get little or no exposure to control ideas at all. This book attempts a rational compromise.[3]

Finally, this book includes topics of recent interest, such as the stochastic approach to robust control, network control, quantum control, and the connections between control and fundamental issues of thermodynamics and information theory. Indeed, the period of writing this book has coincided with a renewed interest in the foundations of control theory.[4] Notions such as controllability, which seemed settled 50 years ago, have been reopened, with sometimes surprising results. And issues such as Maxwell's demon that had befuddled researchers since the dawn of statistical physics 150 years ago have been greatly clarified. While this book focuses on basic theory in control, it also aims to show where new areas fit with more traditional ones and to be a bridge to such recent work.

[3] Åström and Murray (2008) have written an excellent control-theory book for general scientists. They seek a wider audience, including those less comfortable with mathematics. Although full of insights and a broad range of examples, the style, level, and choice of topics are collectively rather different from this book.

[4] The interest can be measured by the recent spate of control-theory articles in *Nature*, *Physical Review Letters*, and *PNAS*. These journals had published little on these topics for many years, until recently.

A word on approach: a little intuition can be worth a lot of formalism, and intuition can be formed through detailed analysis of simple representative examples. That was Einstein's sentiment when he stated that "Everything should be as simple as possible, but not simpler." Typically, I solve the simplest nontrivial example in the text (analytically, if possible) and then elaborate in the problems. There are only a few applications, but there is a thorough treatment of familiar generic examples (e.g., first- and second-order systems, the pendulum). Your job is to transfer acquired intuitions to your own applications!

Figures and Style

The gray-scale figures in this book have been drawn using a minimalist style inspired by the work of Edward Tufte and others (Tufte, 1990, 2001; Frankel and DePace, 2012). Similarly, my writing tries to follow William Strunk's famous advice to "Omit needless words":

> Vigorous writing is concise. A sentence should contain no unnecessary words, a paragraph no unnecessary sentences, for the same reason that a drawing should have no unnecessary lines and a machine no unnecessary parts. This requires not that the writer make all his sentences short, or that he avoid all detail and treat his subjects only in outline, but that every word tell.
>
> (Strunk, W., Jr., 1918)

Although the figure design tends to minimalism, there are hundreds of them: control is a discipline that takes well to illustrations.

Problems

"The problems are an essential part of the text." Such platitudes are often found in textbook prefaces, but here it really is true. I try to demystify the problems as much as possible, while still leaving something for you to do. The goal is that you should be able to know whether you have done the problem correctly. Sometimes, I ask you to "Show that," or to add decimals to an approximate answer, or to reproduce a graph. The main text tries to be "lean but friendly," focusing on the main ideas. Hopefully, this creates a clear narrative, while the problems fill in steps and explore further cases and applications.

Solving problems is the difference between being a tourist in a new land and trying to live there. And who knows? Perhaps you may even decide to stay awhile.

Acknowledgments

I am grateful for the support of many friends and colleagues: Albert Libchaber, Jean-Christophe Géminard, Normand Fortier, and Ray Goldstein first encouraged me to take my interest in control theory seriously. Simon Capelin proposed a book and was incredibly patient; Nicholas Gibbons, along with Róisín Munnelly and Henry Cockburn, saw it to completion. Suckjoon Jun hosted a stay at the Harvard Bauer Center, where the first part of this book was written; Massimiliano Esposito invited me to Luxembourg to finish it off; and Stanislas Leibler invited me for brief but very productive stays at the Simons Center for Systems Biology at the Institute for Advanced Study in Princeton. Thanks to all! I also thank Ben Adcock, Karl Åström, Leslie Ballentine, Jules Bloomenthal, Braden Brinkman, Jean Carlson, Laurent Daudet, Jean-Charles Delvenne, Michel Devoret, Steve Dodge, Max Donelan, Guy Dumont, Alex Fields, Mark Freeman, Barbara Frisken, Harvey Gould, Paul Haljan, Shalev Itzkovitz, Philipp Janert, Malcolm Kennett, John Milton, Elisha Moses, Dugan O'Neill, Rik Pintelon, Max Puelma Touzel, Riccardo Rao, Henrik Sandberg, Pablo Sartori, David Sivak, Philipp Strasberg, Juzar Thingna, Hugo Touchette, Anand Yethiraj, Joel Zylberberg, and the students of Physics 493/881 for suggestions, critiques, and insights. Raphaël Chétrite read the entire manuscript and caught many mistakes. The ones left are my fault! In Canada, NSERC and Simon Fraser University supported my work throughout the writing. I thank Guy Immega for supplying the cover image.

Most of all, I thank Adrienne, qui reste au coeur du tout.