Preface

In the course of the past decade, a great deal of progress has been made in the theory and application of ideas in nonlinear control theory to mechanical systems. The areas of application of control theory for mechanical systems are diverse and challenging, and this constitutes an important factor for the interest in these systems. Such areas of application include robotics and automation, autonomous vehicles in marine, aerospace, and other environments, flight control, problems in nuclear magnetic resonance, micro-electromechanical systems, and fluid mechanics. What is more, the areas of overlap between mechanics and control possess the sort of mathematical elegance that makes them appealing to study, independently of applications.

This book is an outgrowth of our research efforts in the area of mechanics and control, and is a development of separate courses offered by us at our respective institutions, at both the advanced undergraduate and graduate levels. The book reflects our point of view that differential geometric thinking is useful, both in nonlinear control theory and in mechanics. Indeed, the primary emphasis of the book is the exploration of areas of overlap between mechanics and control theory for which differential geometric tools are useful. This area of overlap can be essentially characterized by whether one adopts a Hamiltonian or a Lagrangian view of mechanics. In the former, symplectic or Poisson geometry dominates. On the Lagrangian side, it is less clear what is the most useful geometric framework. We sidestep this to some extent by restricting our attention to the special class of mechanical systems known in the literature as "simple."[1] Simple mechanical systems are characterized by the fact that their Lagrangian is kinetic energy minus potential energy. While it is certainly true that simple mechanical systems are not completely general, it is also true that a very large number of applications, indeed the majority

[1] This terminology appears to originate with Smale [1970]. As the reader might conclude after going through the material in our book, the word "simple" should not be taken to have its colloquial meaning. In some of the literature on mechanics, those systems that we call "simple" are called "natural."

considered in the research literature, fall into the "simple" category. Simple mechanical systems also offer the advantage of providing the useful geometric structure of a Riemannian metric, with its attendant Levi-Civita affine connection. It is this Riemannian geometric, or more generally affine differential geometric, point of view that dominates the presentation in this book.

Intended audience

This book is intended to serve as a reference book for graduate students and for researchers wishing to learn about the affine connection approach to control theory for mechanical systems. The book is also intended to be a textbook for graduate students and undergraduates with a suitable background. Let us say a few words about what a suitable background might be, and then a few words about which portion of the research audience we are aiming for.

Prerequisites. Students who use this book as a text should have a background in analysis, linear algebra, and differential equations. An analysis course, rigorously covering such topics as continuity, differentiability, and convergence, is essential. At Queen's University the book has been used as a text for students with a second-year course in analysis. Linear algebra beyond a first course is also essential. Students will be expected to have a knowledge of abstract vector spaces, linear maps, norms, and inner products. At Queen's a sufficient background is achieved with two courses in linear algebra. A basic course in differential equations may be sufficient, although an advanced course will be very useful. An undergraduate instructor teaching students not having seen this material can be sure to spend a significant amount of time in review using external sources. For graduate students in engineering, our experience is that they will either have had this material, or are capable of filling in the gaps as necessary. At the University of Illinois at Urbana-Champaign the course is offered as a first-year, second-semester graduate course. Students are required to take a prerequisite introductory graduate course on linear control theory. More will be said later in this preface on the use of the book as a text.

A note to researchers. We are certainly aware that a great deal of valuable research is done on control theory for mechanical systems without the aid of the mathematical methodologies advocated by this book. We acknowledge that certain examples of mechanical control systems, and certain general problems in the control theory of mechanical systems, can be treated effectively without the aid of differential geometry. Furthermore, we are sympathetic to the fact that a significant effort is required in order for the reader to become comfortable with the necessary mathematics. However, at some point the unity offered by a differential geometric treatment becomes advantageous and we feel that this is merely a necessary part of the subject, as we see it. It was our objective to write a book about a *class* of mechanical control systems, and not about specific examples. In our approach, examples provide motivation, and they provide a testing ground for general ideas. Nonetheless, we hope that the effectiveness of our ideas, as applied to specific examples, will provide some impetus to the reader unsure of whether they should invest the effort in learning the necessary background. Additionally, one of the features of much of the literature on geometric mechanics and on geometric nonlinear control theory, is that it is written in a precise, mathematical style. We have chosen not to deny this mathematical sophistication. Rather, we have adopted the approach of trying to prepare the uninitiated reader, by providing, under one cover, the necessary background as well as the motivation for what is a demanding mathematical formalism.

Objectives

We state our objectives roughly in descending degree of generality.

Broad aim. Broadly, our objectives were to write a book that could serve as a textbook for instructors with some knowledge of the subject area, and to write a book that would be a reference book, stating some of the important results in the field, and providing a guide to the literature for others. These objectives need not be at odds. Indeed, in terms of a graduate text, the objectives align quite nicely, since students can at the same time learn the basics of the subject, and get a glimpse at some advanced material. For example, at the University of Illinois at Urbana-Champaign this text has been used as an introduction to current research results on nonlinear control of mechanical systems. At the same time, the text has also been used for an advanced undergraduate course at Queen's University. In this capacity the book serves to illustrate to students the value of certain mathematical ideas in the physical sciences, and it serves as a spark to some students to pursue research in related fields.

Control theory for mechanics, and mechanics for control theory. Another objective of the book is to provide a background in (parts of) geometric mechanics for researchers familiar with geometric nonlinear control, and to provide a background in (parts of) geometric nonlinear control for readers familiar with geometric mechanics. This objective is served in two ways. For readers familiar with geometric nonlinear control theory, we describe in some detail the mathematical structure of the physical models we consider. We also clarify the connections between (some of) the tools of the nonlinear control theoretician (e.g., Lyapunov functions, Lie brackets) and (some of) the tools of the geometric mechanic (e.g., Riemannian metrics, nonholonomic constraints). For the reader with a background in geometric mechanics, we provide, in each chapter dealing with matters of control theory, a thorough overview of that area of control theory. It should be noted that the resulting coverage of control theory we provide is quite biased to serving our sometimes rather focused objectives. Therefore, a reader should not feel as if our treatment of nonlinear control is even close to comprehensive. We refer the reader to the texts and monographs [Agrachev and Sachkov 2004, Bloch 2003, Isidori 1995, 1999, Jurdjevic 1997, Khalil 2001, Nijmeijer and van der Schaft 1990, Sastry 1999, Sontag 1998, van der Schaft 1999] for an overview of the subject from multiple points of view.

A unified treatment of the subject. At a more detailed level, the objective of the book is to present a unified treatment of modeling, analysis, and design for mechanical control systems, just as suggested by the title of the book. The unifying feature is our reliance on Riemannian and affine differential geometry to provide the structure in our models. Therefore, a significant portion of the book is devoted to drawing a clear line from physics to differential geometry. We are able to model a significantly large class of systems in our framework, and examples and exercises in each chapter amply illustrate this.

An outline of what is new in this book

It might be of interest to a researcher interested in mechanical control systems to know what material in this book is presented in a novel way or is not present in existing texts on geometry, mechanics, or control. One of the features of the book is that it is divided into three separate parts: Part I, dealing with modeling, Part II, dealing with analysis of mechanical systems and mechanical control systems, and Part III, dealing with design methodologies. In what follows we briefly overview each of these parts.

Part I. Modeling. As a part of modeling, we provide a self-contained review of the mathematical background for the book: linear algebra and differential geometry. In this review, we cover a couple of nonstandard topics that are useful for us later in the book. These topics include the structure of immersed submanifolds (useful for controllability), some background on the character of real-analytic maps, especially contrasted with the behavior of infinitely differentiable maps (useful for understanding certain of the results for which analyticity is required), and the structure of generalized subbundles (useful for controllability and for nonholonomic mechanics). Particularly noteworthy is our rather thorough treatment of distributions. The actual physical modeling is done in a manner that is, as far as we are aware, novel. We present a modeling methodology for mechanical systems that is extremely systematic, and which turns physical data into a differential geometric model. These models form the backbone of the material on analysis and design that follows.

Part II. Analysis. The analysis results we give are for stability, controllability, and perturbation theory. In all cases, we work as much as possible with our intrinsic geometric system models. For stability, we give a coordinate-free notion of exponential stability, and understand the global structure of stability for mechanical systems via Lyapunov analysis. The presentation we give for controllability differs from some standard treatments in that the development of accessibility has as its basis the classic results of Sussmann and Jurdjevic [1972]. We also provide some state-of-the-art local controllability tests, certain of which are not commonly found in texts, and certain of which are, at the time of press, very recent. The coordinate-free approach is carried over to the treatment of averaging and series expansions for mechanical systems. Of particular interest is the use of the intrinsic variation of constants formula for the averaging analysis.

Part III. Design. The geometric formulation of mechanical systems also has real value when it comes to control design. Of the four design chapters, three deal with stabilization, two deal with tracking, and one deals with motion planning. (For those counting, this means that two chapters deal with both stabilization and tracking.) The stabilization results are geometric in flavor, based as they are upon the intrinsic stability and averaging theory developed in the second part of the book. Similarly, the tracking results emphasize the geometry inherent in mechanical control systems. The power of the differential geometric formulation is particularly evident in our consideration of motion planning problems. Merely by understanding the geometry, one is led naturally to a series of seemingly solely mathematical questions. However, the answers to these questions provide simple, explicit motion planning algorithms for systems for which it is not *a priori* clear that such algorithms should exist.

What is not in this book

Although it may not seem commensurate with the length and density of this book, certain choices had to be made concerning what material to present. The need to make such choices was made more pressing in light of our intention to make the book as self-contained as possible. What follows is an incomplete list of topics which we might have included in a manuscript with no length restrictions.

Physics. In the text, we develop a mathematical modeling methodology which we claim is applicable to a large number of physical systems. However, we do not really justify this claim of applicability. We believe this is a serious omission. In our defense we make two points. The first is that Chapters 4 and 5 concerning mathematical modeling already take up a significant proportion of the book, and to develop the physics with an appropriate degree of rigor would have amounted to writing a separate book on mechanics. This leads to the second point of our defense, namely that the physical justification of the applicability of Lagrangian mechanics is taken up in many places, including in the books [Chetaev 1987, Goldstein 1980, Neĭmark and Fufaev 1972, Papastavridis 2002, Pars 1965, Rosenberg 1977, Whittaker 1904]. While none of these books makes the connection between physics and Lagrangian mechanics in the style of our book, perhaps the reader can obtain from them some comfort level with the use of Lagrangian mechanics for modeling physical systems.

Anything Hamiltonian. Many, including the authors, feel that the Hamiltonian setting has a significant role to play in control theory for mechanical systems. In this book, we have focused on a Lagrangian formulation for control theory of mechanical systems. We do this not because we think it superior in any way, but merely because this is where we have done most of our work. The Hamiltonian point of view is treated in Chapter 12 of [Nijmeijer and van der Schaft 1990], and in the recent book of Bloch [2003]. We refer the reader to Section 4.6.6 for additional references.

Complete treatment of symmetry. The subject of symmetry in mechanics is an important one, and is one that we barely scratch the surface of in Chapter 5, and in material related to that chapter. We feel that it is not necessary for us to provide a detailed treatment of symmetry because of the existence of very good books dealing with the subject, including those of Marsden and Ratiu [1999] and Bloch [2003].

Optimal control. For many researchers, particularly some who approach control theory from a certain mathematical point of view, "control theory" is synonymous with "optimal control theory." For such people, our omission of optimal control will be a serious one. The reasons for the omission from the book are exactly space related. However, the good news is that a discussion of optimal control will be included in the supplementary material on the website for the book (the URL is given below). We also refer the reader to the papers [Crouch and Silva Leite 1991, Noakes, Heinzinger, and Paden 1989, Silva Leite, Camarinha, and Crouch 2000] and to the book of Bloch [2003] for a treatment of some problems in optimal control, done using the sorts of tools used in our book.

Advanced control design methodologies. By focusing on control theory for mechanical systems, we have correspondingly focused our treatment of topics in nonlinear control theory. This means that we are not giving systematic treatment to enormous areas of nonlinear control theory. This is particularly true of stabilization theory, where much work has been done. Some standard references for nonlinear control are given in Section 1.4, and for references to additional topics in stabilization we refer the reader to Sections 10.1 and 10.5.

How to use this book as a text

As previously mentioned, various versions of this book have been used as a text, at both the undergraduate and graduate levels. In this section we quickly summarize the content of these courses, and provide some ideas to assist a lecturer who might wish to use this book as a text. We will also have on the webpage for the book (the URL is given below) an expanded version of this section, giving more details, and suggesting some natural variations of material that may make for a coherent course.

The book has been used as a text for an advanced undergraduate course, taught to fourth-year engineering students who are enrolled in a program that is mathematically enriched. There are occasionally mathematics undergraduates, and graduate students from both engineering and mathematics in the course. For this audience, in a twelve week course, one can cover a significant portion of Chapters 3 and 4, along with a collection of topics in control theory. The latter depend largely on the interest of the instructor. The most challenging part of teaching this course is the delivery of the background in differential geometry. This is done, quite effectively, by teaching the material from Chapters 3 and 4 concurrently. In this way, all mathematical concepts are given a fairly immediate physical basis.

A graduate course, taught to engineering students in their first-year, has also been taught from the book. Parts of this course can be sped up as compared to its undergraduate counterpart. It is also true that graduate students in the course will have had a course in linear systems theory, and this will enable the teaching of some deeper control theoretic material. This will be more so, the more familiar the instructor is with the material in the book.

A word of warning is also in order for instructors. The exercises in the book vary greatly in difficulty. In order to not turn students into enemies, instructors would be well-advised to carefully consider the problems they assign. While the occasional difficult problem is good for the constitution, a steady diet of these will probably not make for a good learning experience.

To conclude, let us also mention that we envision complementing and extending this material with internet supplements at the URL for the book:

`http://penelope.mast.queensu.ca/smcs/`

Acknowledgements

It is our great pleasure to thank the various collaborators that have contributed essentially to the material that forms part of this book. Particularly, we acknowledge the contributions of Jorge Cortés, Ron Hirschorn, Naomi Leonard, Kevin Lynch, Sonia Martínez, Richard Murray, David Tyner, and Miloš Žefran. Numerous topics covered in this book are the result of these fruitful collaborations.

We also owe a great debt of gratitude to Richard Murray, who put up with us as graduate students, contributed to our early work, and, most importantly, gave us our initiation to the beautiful subjects of geometric mechanics and mathematical control theory. While students at Caltech, we also had the very good fortune of interacting with Jerry Marsden. Jerry, both in person and through his written work, has provided us and many young academics a model to which to aspire. In the initial stages of development, both Jerry and Richard were also very encouraging, and this was of great value in providing an impetus to continue the work on the book.

We have each used some parts of this book as texts for courses, and we would like to extend our thanks to the students for putting up with the roughness of these early versions of the manuscript, and for finding some of the many small errors that appear in a written work of this size. We also distributed early versions of the manuscript to various graduate students and colleagues, and we received from them a good deal of feedback. In this regard, we particularly wish to thank Ajit Bhand, John Chapman, Peng Cheng, Jorge Cortés, Elsa Hansen, Sonia Martínez, Sina Ober-Bloebaum, Ashoka Polpitiya, Witold Respondek, Jinglai Shen, Stefano Stramigioli, David Tyner, and Miloš Žefran. While many people found errors in the text, we, of course, are responsible for those that remain. On the webpage maintained by the authors can be found errata.

The anonymous reviewers for the book were helpful in providing advice on structure and content. They, along with Achi Dosanjh at Springer, also provided encouragement for the project in the crucial early stages. We would like to thank Manuel de León for his hospitality at the Consejo Superior de Investigaciones Científicas in Madrid during part of the summer of 2003. We are grateful for support from the National Science Foundation in the United States, the Engineering and Physical Sciences Research Council in the United Kingdom, and the Natural Sciences and Engineering Research Council in Canada.

Finally, we acknowledge the use of open source software in the preparation of this book.

Francesco's acknowledgements: I would like to take this opportunity to truly thank my wife, Lily, for her patience and support during all the long hours this project required; her loving presence made all the difference. I sincerely thank my parents, Carla and Aurelio, and my siblings, Valentina and Federico, for all their love and encouragement. Although we were geographically separated, they have always been there for me. Finally, I extend a world of thanks to my second set of parents, Martha and Juan, and to my brother-in-law, Carlos, for sharing so many memorable moments during the writing of this book.

Andrew's acknowledgements: I extend my gratitude to Laurie for putting up with the long hours required to write this book. I hope it proves to be worthwhile. While this book was in preparation, I experienced two singular events. In March 2001, Laurie and I gave birth to our son Gabriel, who has become my very best friend. His presence has provided me with a much needed distraction of pure happiness. In February 2003 my mother Elizabeth passed away. It is not possible to express the positive influence she has had, and will continue to have, on my life. She will be greatly missed, particularly by my sister Cathryn and myself, and it is to her that I primarily dedicate my share in this book.

FB
Urbana, IL
United States

ADL
Kingston, ON
Canada