


Undergraduate Texts in Mathematics

Jack Macki
Aaron Strauss

**Introduction
to Optimal
Control Theory**



Springer-Verlag New York • Heidelberg • Berlin

TP13
M2

8262332

Jack Macki
Aaron Strauss

Introduction to Optimal Control Theory



E8262332



Springer-Verlag
New York Heidelberg Berlin

1182

Jack Macki
Department of Mathematics
The University of Alberta
Edmonton, Canada T6G 2G1

Aaron Strauss
formerly of
Division of Mathematical and
Physical Sciences and Engineering
University of Maryland
College Park, MD 20742
USA

Editorial Board

P. R. Halmos
Managing Editor
Department of Mathematics
Indiana University
Bloomington, IN 47401
USA

F. W. Gehring
Department of Mathematics
University of Michigan
Ann Arbor, MI 48108
USA

C. C. Moore
Department of Mathematics
University of California
Berkeley, CA 94720
USA

AMS Subject Classification: 49-01

With 68 Figures.

Library of Congress Cataloging in Publication Data

Macki, Jack.

Introduction to optimal control theory.

(Undergraduate texts in mathematics)

Bibliography: p.

Includes index.

1. Control theory. I. Strauss, Aaron.

II. Title. III. Series.

QA402.3.M317 629.8'312 81-8889

AACR2

© 1982 by Springer-Verlag New York Inc.

All rights reserved. No part of this book may be translated or reproduced in any form without written permission from Springer-Verlag, 175 Fifth Avenue, New York, New York 10010, U.S.A.

Printed in the United States of America.

9 8 7 6 5 4 3 2 1

ISBN 0-387-90624-X Springer-Verlag New York Heidelberg Berlin
ISBN 3-540-90624-X Springer-Verlag Berlin Heidelberg New York

Undergraduate Texts in Mathematics

Editors

F. W. Gehring

P. R. Halmos

Advisory Board

C. DePrima

I. Herstein

Dedicated to the memory of Aaron Strauss
(1940–1977)

Preface

This monograph is an introduction to optimal control theory for systems governed by vector ordinary differential equations. It is not intended as a state-of-the-art handbook for researchers. We have tried to keep two types of reader in mind: (1) mathematicians, graduate students, and advanced undergraduates in mathematics who want a concise introduction to a field which contains nontrivial interesting applications of mathematics (for example, weak convergence, convexity, and the theory of ordinary differential equations); (2) economists, applied scientists, and engineers who want to understand some of the mathematical foundations of optimal control theory.

In general, we have emphasized motivation and explanation, avoiding the "definition-axiom-theorem-proof" approach. We make use of a large number of examples, especially one simple canonical example which we carry through the entire book. In proving theorems, we often just prove the simplest case, then state the more general results which can be proved. Many of the more difficult topics are discussed in the "Notes" sections at the end of chapters and several major proofs are in the Appendices. We feel that a solid understanding of basic facts is best attained by at first avoiding excessive generality.

We have not tried to give an exhaustive list of references, preferring to refer the reader to existing books or papers with extensive bibliographies. References are given by author's name and the year of publication, e.g., Waltman [1974].

Prerequisites for reading this monograph are basic courses in ordinary differential equations, linear algebra, and modern advanced calculus (including some Lebesgue integration). Some functional analysis is used, but the proofs involved may be treated as optional. We have summarized the


relevant facts from these areas in an Appendix. We also give references in this Appendix to standard texts in these areas.

We would like to express our appreciation to: Professor Jim Yorke of the University of Maryland for providing several important and original proofs to simplify the presentation of difficult material; Dr. Stephen Lewis of the University of Alberta for providing several interesting examples from Economics; Ms. Peggy Gendron of the University of Minnesota and June Talpash and Laura Thompson of Edmonton, Alberta for their excellent typing work; the universities (and, ultimately, the relevant taxpayers) of Alberta, Maryland, and Minnesota – the first two for their direct financial support, and the last for providing facilities for J.W.M. while on sabbatical; The National Research Council of Canada, for its continuing support of J.W.M.

Edmonton, Alberta
August, 1980

Jack W. Macki

List of Symbols



$\mathcal{B}(\mathbf{x}; \alpha)$	the open ball about \mathbf{x} of radius α , $\{\mathbf{y} \mathbf{y} - \mathbf{x} < \alpha\}$
$C, C[\mathbf{u}(\cdot)]$	cost function
$C(\mathbf{x}_0)$	cost function evaluated along the optimal control and response from \mathbf{x}_0
C^1	class of functions having continuous first partial derivatives
C^∞	functions having continuous partials of every order
C_0^∞	functions from C^∞ which have compact support
$\text{co}(\Omega)$	convex hull of Ω
\mathcal{C}	controllable set
$\mathcal{C}(t)$	controllable set at time t
\mathcal{C}_{BB}	controllable set using bang-bang controls
$\mathcal{C}_{\text{BB}}(t)$	controllable set at time t using bang-bang controls
$\mathcal{C}_{\text{BBPC}}$	controllable set using bang-bang piecewise constant controls
$d(\mathbf{x}, P)$	$\inf \{ \mathbf{x} - \mathbf{y} : \mathbf{y} \in P\}$
∂S	boundary of the set S
f^0	$C[\mathbf{u}(\cdot)] = \int_0^t f^0(s, \mathbf{x}[s], \mathbf{u}(s)) ds$
\mathbf{f}_x	matrix of partials $\partial f^i / \partial x^j$
$\hat{\mathbf{f}}$	$(f^0, \mathbf{f}^T)^T$; the extended velocity vector
$\text{grad}_x H$	gradient = $\left(\frac{\partial H}{\partial x^1}, \dots, \frac{\partial H}{\partial x^n} \right)$

$h(P, Q)$	Hausdorff metric = $\inf \{ \varepsilon : P \subset N(Q, \varepsilon) \text{ and } Q \subset N(P, \varepsilon) \}$
$H, H(\hat{\mathbf{w}}, \mathbf{x}, \mathbf{u})$	$\langle \hat{\mathbf{w}}, \hat{\mathbf{f}}(\mathbf{x}, \mathbf{u}) \rangle$
Int S	the interior of the set S
$k(\tau)$	$\left\{ \sum_{i=1}^P c^i [\hat{\mathbf{f}}(x_*[\tau], \mathbf{v}_i) - \hat{\mathbf{f}}(x_*[\tau], \mathbf{u}_*(\tau))] c^i \geq 0, \mathbf{v}_i \in \Psi, P \in \mathcal{N} \right\}$; here $(x_*[\cdot], y_*[\cdot])$ is optimal and Ψ is the range set for admissible controls
$K(t; \mathbf{x}_0)$	reachable set at time t
$K_{\text{BB}}(t; \mathbf{x}_0)$	reachable set at time t using bang-bang controls
$\mathcal{K}(t)$	$\left\{ \sum_{i=1}^P c^i Y(t, \tau_i) \hat{\mathbf{z}}_i \mid c_i \geq 0, \hat{\mathbf{z}}_i \in k(\tau_i) \right\}$, where $Y(t, \tau_i)$ is the fundamental matrix for (Lin) satisfying $Y(\tau_i, \tau_i) = I$
$\bar{\mathcal{K}}(t_1)$	$\{ \hat{\mathbf{a}} + \beta \hat{\mathbf{b}} \mid \beta \leq \beta_0, \hat{\mathbf{a}} \in \mathcal{K}(t_1), \hat{\mathbf{b}} = \hat{\mathbf{f}}(\mathbf{x}_*[t_1], \mathbf{u}_*(t_1)) \}$
(L)	$\dot{\mathbf{x}} = A(t)\mathbf{x} + B(t)\mathbf{u} + \mathbf{c}(t)$
(LA)	$\dot{\mathbf{x}} = A\mathbf{x} + B\mathbf{u}$
m	dimension of control vectors $\mathbf{u}(t)$
$M(\hat{\mathbf{w}}, \mathbf{x})$	$\sup \{ H(\hat{\mathbf{w}}, \mathbf{x}, \mathbf{v}) : \mathbf{v} \in \Omega \}$
$M \equiv (B, AB, \dots, A^{n-1}B)$	the controllability matrix for (LA)
$N(P, \varepsilon)$	$\{ \mathbf{x} : d(\mathbf{x}, P) < \varepsilon \}$
$o(\mathbf{x})$	stands for any function $h(\mathbf{x})$ such that $\lim_{\mathbf{x} \rightarrow 0} \frac{h(\mathbf{x})}{ \mathbf{x} } = 0$
$Q^+(t, \mathbf{x})$	$\{ (y^0, \mathbf{y}) \mid \exists \mathbf{v} \in \Omega, \mathbf{y} = \mathbf{f}(t, \mathbf{x}, \mathbf{v}), y_0 \geq f^0(t, \mathbf{x}, \mathbf{v}) \}$
RC	reachable cone, $\bigcup_{t > t_0} (t, K(t, \mathbf{x}_0))$
R^n	Euclidean n -dimensional space
sgn α	$\alpha/ \alpha $ provided $\alpha \neq 0$
$\mathcal{T}(t)$	target state
\mathcal{U}_{BB}	class of functions in \mathcal{U}_m for which $ u^i(t) = 1$
\mathcal{U}_m	$\bigcup_{t_1 > t_0} \mathcal{U}_m(t_0, t_1)$
$\mathcal{U}_m(t_0, t_1)$	class of measurable functions from $[t_0, t_1]$ to Ω
\mathcal{U}_{PC}	class of piecewise constant functions in \mathcal{U}_m
\mathcal{U}_{PS}	class of piecewise smooth functions in \mathcal{U}_m
\mathcal{U}_r	class of piecewise constant functions in \mathcal{U}_m with at most r discontinuities

\mathcal{U}_λ	class of functions in \mathcal{U}_m having Lipschitz constant λ
$\mathcal{V}_m(t_0, t_1)$	class of measurable functions from $[t_0, t_1]$ to a given bounded set $\Psi \subset R^m$
\mathcal{V}_m	$\bigcup_{t_1 > t_0} \mathcal{V}_m(t_0, t_1)$
$\mathbf{x}(t; t_0, \mathbf{x}_0, \mathbf{u}(\cdot))$	solution of relevant differential equation through \mathbf{x}_0 at time t_0 corresponding to $\mathbf{u}(\cdot)$; the <i>state</i> vector
x^i	i^{th} component of \mathbf{x}
\mathbf{x}^T	transpose of \mathbf{x}
$\hat{\mathbf{x}}$	$(x^0, \mathbf{x}) \in R^{n+1}$; the extended state vector
$\langle \mathbf{x}, \mathbf{y} \rangle$	$\sum_i x^i y^i$
Δ	class of successful controls: they steer the initial state to the target
χ_Q	characteristic function of a set Q , i.e., $\chi = +1$ on Q , 0 on the complement of Q
Ω	the unit cube in R^m

Contents

List of Symbols	xi
Chapter I	
Introduction and Motivation	1
1 Basic Concepts	1
2 Mathematical Formulation of the Control Problem	4
3 Controllability	6
4 Optimal Control	9
5 The Rocket Car	10
Exercises	15
Notes	20
Chapter II	
Controllability	24
1 Introduction: Some Simple General Results	24
2 The Linear Case	28
3 Controllability for Nonlinear <u>Autonomous</u> Systems	38
4 Special Controls	44
Exercises	48
Appendix: Proof of the Bang–Bang Principle	50
Chapter III	
Linear Autonomous Time-Optimal Control Problems	57
1 Introduction: Summary of Results	57
2 The Existence of a Time-Optimal Control; Extremal Controls; the Bang–Bang Principle	60

3 Normality and the Uniqueness of the Optimal Control	65
4 Applications	74
5 The Converse of the Maximum Principle	77
6 Extensions to More General Problems	79
Exercises	80
Chapter IV	
Existence Theorems for Optimal Control Problems	82
1 Introduction	82
2 Three Discouraging Examples. An Outline of the Basic Approach to Existence Proofs	83
3 Existence for Special Control Classes	88
4 Existence Theorems under Convexity Assumptions	91
5 Existence for Systems Linear in the State	97
6 Applications	98
Exercises	100
Notes	102
Chapter V	
Necessary Conditions for Optimal Controls—The Pontryagin Maximum Principle	103
1 Introduction	103
2 The Pontryagin Maximum Principle for Autonomous Systems	104
3 Applying the Maximum Principle	111
4 A Dynamic Programming Approach to the Proof of the Maximum Principle	118
5 The PMP for More Complicated Problems	124
Exercises	128
Appendix to Chapter V—A Proof of the Pontryagin Maximum Principle	134
Mathematical Appendix	147
Bibliography	160
Index	163

Chapter I

Introduction and Motivation

1. Basic Concepts

In control theory, one is interested in governing the *state* of a *system* by using *controls*. The best way to understand these three concepts is through examples.

EXAMPLE I (A National Economy). The economy of a typical capitalistic nation is a *system* made up in part of the population (as consumers and as producers), companies, material goods, production facilities, cash and credit available, and so on. The *state* of the system can be thought of as a massive collection of data: wages and salaries, profits, losses, sales of goods and services, investment, unemployment, welfare costs, the inflation rate, gold and currency holdings, and foreign trade. The federal government can influence the state of this system by using several *controls*, notably the prime interest rate, taxation policy, and persuasion regarding wage and price settlements.

EXAMPLE II (Water Storage and Supply). As early as the third century B.C., systems similar to that sketched in Figure 1 were being used in water storage tanks. As the water level rises, the float will restrict the inlet flow; all inlet flow will cease when the water reaches a certain height. If water is withdrawn from the outlet at a certain rate then the float will tend to adjust the inlet flow so as to maintain the water height in the tank. One can think of the water in the tank along with the float, inlet, and outlet, as a *system*. The *control* is the position of the float. The *state* at any instant is a vector, consisting of the height of the water in the tank, the inlet rate of flow and the outlet rate of flow. In this example, the state of the system




Figure 1

(rather than an external observer) automatically sets the control (position of the float) – this is an example of a *feedback* control system – the state is “fed back” to the control mechanism, which adjusts without outside influence.

We have chosen one example from economics and one from civil engineering. We could just as well have chosen examples from biology, economics, space flight, or several other fields, because the concepts of system, state, and control are so general. In the exercises at the end of the chapter we have given several more examples.

The essence of these examples is that in control theory we have a system and we try to influence the state of the system through controls. The *dynamics* of the system, that is, the manner in which the state changes under the influence of the controls, can be very complicated in real-world examples. In the case of a national economy, the dynamics is still a matter of considerable research. Of course, there are many general principles for a national economy – for example, raising the prime rate (a control) generally increases unemployment – but a detailed, accurate picture of the dynamics of a national economy is very difficult. On the other hand, the dynamics of the water storage system is relatively easy to describe. We won’t do it here, since we are going to deal with an even simpler example shortly.

There are two remaining concepts to be described, namely the *constraints* on our controls, and the *objective* or *target* state(s) for our system. For a national economy, there are several obvious constraints on our controls, for example, taxation cannot be too excessive and the prime rate cannot be negative. There are also objective or target states – ideally a government wants a state of the economy with full employment, an inflation rate of 0%, low interest rates, and low taxes. In fact, they may have to settle for a realistic target state with an unemployment rate less than 8%, inflation less than 10%, moderate interest rates, and realistic tax rates. *Any* state with these properties would do, so there are many target states. In fact, the set of target states might vary with time, reflecting political and social changes.

In the water tank example, the constraint on the position of the stopper is that it always floats at a fixed distance above the water level; also the velocity with which the stopper can move is tied directly to the rate of change of the water level height. The objective might be a pre-set water height.

EXAMPLE 1 (The Rocket Car). This example will be used throughout this monograph to motivate and illustrate concepts and results. The car runs on rails on the level, has a mass of one, and is equipped with two rocket engines, one on each end (Figure 2). The problem is to move the car from




Figure 2

any given location to a fixed pre-assigned destination. For simplicity, we place the destination at the origin and denote the position of the center of the car by $p(t)$. If the car is at a position p_0 at time $t = 0$, with velocity v_0 , we want to fire the two engines according to some recipe (pattern, program) which will have us arrive at $p = 0$ at rest (with velocity zero) at some instant $t_1 > 0$. We can take as our *system* the car plus its track; as the *state* we take the two-vector $\mathbf{x}(t) = (p(t), \dot{p}(t))$; the initial state (p_0, v_0) is assumed given. The physical reason for using a two-vector for the state is simple – we want to know where we are and how fast we are going. Our *target* state is $(0, 0)$. A *control* $u(t)$ is a real-valued function, representing the force on the car due to firing either engine at time t . If we fire the right engine at time t^* , we will say the force is negative, if we fire the left engine we take the force positive.




Figure 3 (Moving to the Right) The Force Is to the Left When $u(t^*) < 0$.

Then the *dynamics* of our system is given by Newton's law $F = ma$, which can be written as $\ddot{p}(t) = u(t)$. This has the natural vector form

$$\mathbf{x}(t) = \begin{bmatrix} p(t) \\ \dot{p}(t) \end{bmatrix}, \quad \dot{\mathbf{x}}(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \mathbf{x}(t) + u(t) \begin{bmatrix} 0 \\ 1 \end{bmatrix}.$$

There are *constraints* on the magnitude of $u(t)$, based on the size of the rocket motors and the amount of acceleration stress allowed on the car.

A mathematically reasonable assumption is that $u(t)$ is measurable and bounded, and we take our constraint to be $|u(t)| \leq 1$ for simplicity. Since measurable functions can be quite pathological, we will often use classes that are physically more reasonable, e.g., piecewise constant controls.

A given control function $u(t)$ is a recipe for firing our engines. For example

$$u(t) = \begin{cases} +1, & 0 \leq t \leq 1; \\ -\frac{1}{2}, & 1 < t \leq 3; \end{cases}$$

tells us to fire the left engine at full force for one unit of time, then fire the right engine at half force for two units of time.

If (p_0, v_0) is the position of the car at $t=0$, we can integrate the differential equation twice and then integrate by parts to get:

$$p(t) = p_0 + v_0 t + \int_0^t (t-r)u(r) dr, \quad \dot{p}(t) = v_0 + \int_0^t u(r) dr.$$

Thus each choice of control, $u(\cdot)$, generates a *response* $\mathbf{x}[t] \equiv \mathbf{x}(t; \mathbf{x}_0, u(\cdot))$. (We delete reference to the initial time t_0 , since we always take it to be zero for simplicity. For systems not explicitly containing t , this is in fact not a restriction.) We use $u(\cdot)$ to refer to the function $u(t)$, on its domain of definition, as an entity. If the response $\mathbf{x}(t; \mathbf{x}_0, u(\cdot))$ reaches the target $(0, 0)$ at some $t_1 > 0$, then $u(\cdot)$ is a *successful control*. There might be no such control or many. When there are several successful controls, the choice of one over the other may be dictated by practicality, and/or by a *cost* or *performance criterion*. For example, later on we shall consider the criteria: (1) least time, (2) least energy expended, (3) least fuel expended. Our control problem will then become an *Optimal Control Problem*.

2. Mathematical Formulation of the Control Problem

We now give a precise mathematical formulation of the type of control problem we will be discussing. Let m, n be natural numbers, and let R stand for the real numbers. If \mathbf{x}, \mathbf{y} are column vectors in R^n , we denote their i^{th} components by x^i, y^i respectively. We define \mathbf{x}^T to be the transpose of \mathbf{x} , and introduce a dot product and two norms:

$$\langle \mathbf{x}, \mathbf{y} \rangle = \mathbf{x}^T \mathbf{y} = \sum_{i=1}^n x^i y^i,$$

$$|\mathbf{x}| = \sum_{i=1}^n |x^i|, \quad \|\mathbf{x}\| = \langle \mathbf{x}, \mathbf{x} \rangle^{1/2}.$$

If we need to square a scalar-valued function $\phi(t)$, we will write $[\phi(t)]^2$, while $x^2(t)$ will denote the second component of the vector-valued function $\mathbf{x}(t)$ – in context the distinction will always be obvious. Let Ω denote the unit cube in R^m , i.e.,

$$\Omega = \{\mathbf{c} \mid \mathbf{c} \in R^m, |c^i| \leq 1, i = 1, 2, \dots, m\}.$$

For $t_1 \geq 0$, define

$$\mathcal{U}_m[0, t_1] = \{\mathbf{u}(\cdot) \mid \mathbf{u}(t) \in \Omega \text{ and } \mathbf{u}(\cdot) \text{ measurable on } [0, t_1]\},$$

$\mathcal{U}_m = \bigcup_{t_1 > 0} \mathcal{U}_m[0, t_1]$. Unless explicitly stated otherwise, our controls $\mathbf{u}(\cdot)$ will always be assumed to belong to \mathcal{U}_m . This mildly cumbersome definition of our admissible controls allows each control $\mathbf{u}(\cdot)$ to have its own corresponding interval of definition $[0, t_1(\mathbf{u})]$.

We assume that for each $t \geq 0$ we are given a target set $\mathcal{T}(t) \subset R^n$ where $\mathcal{T}(t)$ is a closed set. For most of this monograph we will take $\mathcal{T}(t) \equiv 0 \in R^n$ for simplicity. Nevertheless, general target sets are important, as we mentioned in the example of a national economy.

We assume that the dynamics of the system, that is, the evolution of the state $\mathbf{x}(t)$ under a given control $\mathbf{u}(t)$, is determined by a vector ordinary differential equation:

$$(1) \quad \dot{\mathbf{x}}(t) = \mathbf{f}(t, \mathbf{x}(t), \mathbf{u}(t)), \quad \mathbf{x}(t_0) = \mathbf{x}_0.$$

We will always assume that $\mathbf{f}(t, \mathbf{x}, \mathbf{u})$, $\partial f^i / \partial x^j$, $\partial f^i / \partial u^k$ are all continuous ($i, j = 1, \dots, n; k = 1, \dots, m$) on $[0, \infty) \times R^n \times R^m$, although most results are valid under weaker conditions. This assumption guarantees local existence and uniqueness of the solution of (1) for a given $\mathbf{u}(\cdot) \in \mathcal{U}_m$. Because $\mathbf{u}(\cdot)$ is only assumed measurable and bounded, the right side of the equation $\dot{\mathbf{x}} = \mathbf{f}(t, \mathbf{x}, \mathbf{u}(t))$ is continuous in \mathbf{x} but only measurable and bounded in t for each \mathbf{x} . Therefore, solutions are understood to be absolutely continuous functions that satisfy (1) almost everywhere. The solution of (1) for a given $\mathbf{u}(\cdot)$ will be called the *response* to $\mathbf{u}(\cdot)$; we denote it by $\mathbf{x}[t] \equiv \mathbf{x}(t; \mathbf{x}_0, \mathbf{u}(\cdot))$. The *control problem* is to determine those \mathbf{x}_0 and $\mathbf{u}(\cdot) \in \mathcal{U}_m$ such that the associated response satisfies $\mathbf{x}[t_1] \in \mathcal{T}(t_1)$ for some $t_1 > 0$; we then say that *the control $\mathbf{u}(\cdot)$ steers \mathbf{x}_0 to the target*.

If the control $\mathbf{u}(\cdot)$ is defined on $[0, t_1)$ ($t_1 \leq +\infty$), it is not assumed that the corresponding response extends to $[0, t_1)$; a given response $\mathbf{x}(t; \mathbf{x}_0, \mathbf{u}(\cdot))$ may only exist on some subinterval of $[0, t_1)$. For example, consider the scalar problem $\dot{x} = x^2 + u$, $x_0 = 1$. For $u_0(t) \equiv 0$ ($t_1 = +\infty$), the response is $x(t; 1, u_0(\cdot)) = 1/(1-t)$, which only exists on $[0, 1)$. For linear equations,

$$(2) \quad \dot{\mathbf{x}}(t) = \mathbf{A}(t)\mathbf{x}(t) + \mathbf{B}(t)\mathbf{u}(t),$$

with $\mathbf{A}(t)$ and $\mathbf{B}(t)$ continuous on $[0, t_1)$, solutions always extend to $[0, t_1)$.

Thus, our general control problem consists of a class of admissible controls \mathcal{U}_m , a vector differential equation (1) describing the *dynamics* of our system, and a family of *target* sets $\mathcal{T}(t)$. One basic problem is to describe