

International Series in Operations Research & Management Science

Volume 196

Series Editor

Frederick S. Hillier
Stanford University, CA, USA

Special Editorial Consultant

Camille C. Price
Stephen F. Austin State University, TX, USA

For further volumes:
<http://www.springer.com/series/6161>

Robert J. Vanderbei

Linear Programming

Foundations and Extensions

Fourth Edition



Springer

Robert J. Vanderbei
Department of Operations Research
and Financial Engineering
Princeton University
Princeton, New Jersey, USA

ISSN 0884-8289
ISBN 978-1-4614-7629-0 ISBN 978-1-4614-7630-6 (eBook)
DOI 10.1007/978-1-4614-7630-6
Springer New York Heidelberg Dordrecht London

Library of Congress Control Number: 2013939593

© Springer Science+Business Media New York 2014

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

*To Krisadee,
Marisa and Diana*

Preface

This book is about constrained optimization. It begins with a thorough treatment of linear programming and proceeds to convex analysis, network flows, integer programming, quadratic programming, and convex optimization. Along the way, dynamic programming and the linear complementarity problem are touched on as well.

The book aims to be a first introduction to the subject. Specific examples and concrete algorithms precede more abstract topics. Nevertheless, topics covered are developed in some depth, a large number of numerical examples are worked out in detail, and many recent topics are included, most notably interior-point methods. The exercises at the end of each chapter both illustrate the theory and, in some cases, extend it.

Prerequisites. The book is divided into four parts. The first two parts assume a background only in linear algebra. For the last two parts, some knowledge of multivariate calculus is necessary. In particular, the student should know how to use Lagrange multipliers to solve simple calculus problems in 2 and 3 dimensions.

Associated software. It is good to be able to solve small problems by hand, but the problems one encounters in practice are large, requiring a computer for their solution. Therefore, to fully appreciate the subject, one needs to solve large (practical) problems on a computer. An important feature of this book is that it comes with software implementing the major algorithms described herein. At the time of writing, software for the following five algorithms is available:

- The two-phase simplex method as shown in Figure 6.1.
- The self-dual simplex method as shown in Figure 7.1.
- The path-following method as shown in Figure 18.1.
- The homogeneous self-dual method as shown in Figure 22.1.
- The long-step homogeneous self-dual method as described in Exercise 22.4.

The programs that implement these algorithms are written in C and can be easily compiled on most hardware platforms. Students/instructors are encouraged to install and compile these programs on their local hardware. Great pains have been taken to make the source code for these programs readable (see Appendix A). In particular, the names of the variables in the programs are consistent with the notation of this book.

There are two ways to run these programs. The first is to prepare the input in a standard computer-file format, called MPS format, and to run the program using such a file as input. The advantage of this input format is that there is an archive of problems stored in this format, called the NETLIB suite, that one can download and use immediately (a link to the NETLIB suite can be found at the web site mentioned below). But, this format is somewhat archaic and, in particular, it is not easy to create these files by hand. Therefore, the programs can also be run from within a problem modeling system called AMPL. AMPL allows one to describe mathematical programming problems using an easy to read, yet concise, algebraic notation. To run the programs within AMPL, one simply tells AMPL the name of the solver-program before asking that a problem be solved. The text that describes AMPL, Fourer et al. (1993) makes an excellent companion to this book. It includes a discussion of many practical linear programming problems. It also has lots of exercises to hone the modeling skills of the student.

Several interesting computer projects can be suggested. Here are a few suggestions regarding the simplex codes:

- Incorporate the partial pricing strategy (see Section 8.7) into the two-phase simplex method and compare it with full pricing.
- Incorporate the steepest-edge pivot rule (see Section 8.8) into the two-phase simplex method and compare it with the largest-coefficient rule.
- Modify the code for either variant of the simplex method so that it can treat bounds and ranges implicitly (see Chapter 9), and compare the performance with the explicit treatment of the supplied codes.
- Implement a “warm-start” capability so that the sensitivity analyses discussed in Chapter 7 can be done.
- Extend the simplex codes to be able to handle integer programming problems using the branch-and-bound method described in Chapter 23.

As for the interior-point codes, one could try some of the following projects:

- Modify the code for the path-following algorithm so that it implements the affine-scaling method (see Chapter 21), and then compare the two methods.
- Modify the code for the path-following method so that it can treat bounds and ranges implicitly (see Section 20.3), and compare the performance against the explicit treatment in the given code.
- Modify the code for the path-following method to implement the higher-order method described in Exercise 18.5. Compare.
- Extend the path-following code to solve quadratic programming problems using the algorithm shown in Figure 24.3.
- Further extend the code so that it can solve convex optimization problems using the algorithm shown in Figure 25.2.

And, perhaps the most interesting project of all:

- Compare the simplex codes against the interior-point code and decide for yourself which algorithm is better on specific families of problems.

The software implementing the various algorithms was developed using consistent data structures and so making fair comparisons should be straightforward. The software can be downloaded from the following web site:

<http://www.princeton.edu/~rvdb/LPbook/>

If, in the future, further codes relating to this text are developed (for example, a self-dual network simplex code), they will be made available through this web site.

Features. Here are some other features that distinguish this book from others:

- The development of the simplex method leads to Dantzig's parametric self-dual method. A randomized variant of this method is shown to be immune to the travails of degeneracy.
- The book gives a balanced treatment to both the traditional simplex method and the newer interior-point methods. The notation and analysis is developed to be consistent across the methods. As a result, the self-dual simplex method emerges as the variant of the simplex method with most connections to interior-point methods.
- From the beginning and consistently throughout the book, linear programming problems are formulated in *symmetric form*. By highlighting symmetry throughout, it is hoped that the reader will more fully understand and appreciate duality theory.
- By slightly changing the right-hand side in the Klee–Minty problem, we are able to write down an explicit dictionary for each vertex of the Klee–Minty problem and thereby uncover (as a homework problem) a simple, elegant argument why the Klee–Minty problem requires $2^n - 1$ pivots to solve.
- The chapter on regression includes an analysis of the expected number of pivots required by the self-dual variant of the simplex method. This analysis is supported by an empirical study.
- There is an extensive treatment of modern interior-point methods, including the primal–dual method, the affine-scaling method, and the self-dual path-following method.
- In addition to the traditional applications, which come mostly from business and economics, the book features other important applications such as the optimal design of truss-like structures and L^1 -regression.

Exercises on the Web. There is always a need for fresh exercises. Hence, I have created and plan to maintain a growing archive of exercises specifically created for use in conjunction with this book. This archive is accessible from the book's web site:

<http://www.princeton.edu/~rvdb/LPbook/>

The problems in the archive are arranged according to the chapters of this book and use notation consistent with that developed herein.

Advice on solving the exercises. Some problems are routine while others are fairly challenging. Answers to some of the problems are given at the back of the book.

In general, the advice given to me by Leonard Gross (when I was a student) should help even on the hard problems: *follow your nose*.

Audience. This book evolved from lecture notes developed for my introductory graduate course in linear programming as well as my upper-level undergraduate course. A reasonable undergraduate syllabus would cover essentially all of Part 1 (Simplex Method and Duality), the first two chapters of Part 2 (Network Flows and Applications), and the first chapter of Part 4 (Integer Programming). At the graduate level, the syllabus should depend on the preparation of the students. For a well-prepared class, one could cover the material in Parts 1 and 2 fairly quickly and then spend more time on Parts 3 (Interior-Point Methods) and 4 (Extensions).

Dependencies. In general, Parts 2 and 3 are completely independent of each other. Both depend, however, on the material in Part 1. The first Chapter in Part 4 (Integer Programming) depends only on material from Part 1, whereas the remaining chapters build on Part 3 material.

Acknowledgments. My interest in linear programming was sparked by Robert Garfinkel when we shared an office at Bell Labs. I would like to thank him for his constant encouragement, advice, and support. This book benefited greatly from the thoughtful comments and suggestions of David Bernstein and Michael Todd. I would also like to thank the following colleagues for their help: Ronny Ben-Tal, Leslie Hall, Yoshi Ikura, Victor Klee, Irvin Lustig, Avi Mandelbaum, Marc Meke-ton, Narcis Nabona, James Orlin, Andrzej Ruszczyński, and Henry Wolkowicz. I would like to thank Gary Folven at Kluwer and Fred Hillier, the series editor, for encouraging me to undertake this project. I would like to thank my students for finding many typos and occasionally more serious errors: John Gilmartin, Jacinta Warnie, Stephen Woolbert, Lucia Wu, and Bing Yang. My thanks to Erhan Çinlar for the many times he offered advice on questions of style. I hope this book reflects positively on his advice. Finally, I would like to acknowledge the support of the National Science Foundation and the Air Force Office of Scientific Research for supporting me while writing this book. In a time of declining resources, I am especially grateful for their support.

Princeton, NJ, USA

Robert J. Vanderbei

Preface to 2nd Edition

For the 2nd edition, many new exercises have been added. Also I have worked hard to develop online tools to aid in learning the simplex method and duality theory. These online tools can be found on the book's web page:

<http://www.princeton.edu/~rvdb/LPbook/>

and are mentioned at appropriate places in the text. Besides the learning tools, I have created several online exercises. These exercises use randomly generated problems and therefore represent a virtually unlimited collection of “routine” exercises that can be used to test basic understanding. Pointers to these online exercises are included in the exercises sections at appropriate points.

Some other notable changes include:

- The chapter on network flows has been completely rewritten. Hopefully, the new version is an improvement on the original.
- Two different fonts are now used to distinguish between the set of basic indices and the basis matrix.
- The first edition placed great emphasis on the symmetry between the primal and the dual (the negative transpose property). The second edition carries this further with a discussion of the relationship between the basic and nonbasic matrices B and N as they appear in the primal and in the dual. We show that, even though these matrices differ (they even have different dimensions), $B^{-1}N$ in the dual is the negative transpose of the corresponding matrix in the primal.
- In the chapters devoted to the simplex method in matrix notation, the collection of variables $z_1, z_2, \dots, z_n, y_1, y_2, \dots, y_m$ was replaced, in the first edition, with the single array of variables y_1, y_2, \dots, y_{n+m} . This caused great confusion as the variable y_i in the original notation was changed to y_{n+i} in the new notation. For the second edition, I have changed the notation for the single array to z_1, z_2, \dots, z_{n+m} .
- A number of figures have been added to the chapters on convex analysis and on network flow problems.
- The algorithm referred to as the primal–dual simplex method in the first edition has been renamed the parametric self-dual simplex method in accordance with prior standard usage.

- The last chapter, on convex optimization, has been extended with a discussion of merit functions and their use in shortening steps to make some otherwise nonconvergent problems converge.

Acknowledgments. Many readers have sent corrections and suggestions for improvement. Many of the corrections were incorporated into earlier reprintings. Only those that affected pagination were accrued to this new edition. Even though I cannot now remember everyone who wrote, I am grateful to them all. Some sent comments that had significant impact. They were Hande Benson, Eric Denardo, Sudhakar Mandapati, Michael Overton, and Jos Sturm.

Princeton, NJ, USA

Robert J. Vanderbei

Preface to 3rd Edition

It has been almost 7 years since the 2nd edition appeared and the publisher is itching for me to finish a new edition. The previous edition had very few typos. I have fixed them all! Of course, I've also added some new material and who knows how many new typos I've introduced. The most significant new material is contained in a new chapter on financial applications, which discusses a linear programming variant of the portfolio selection problem and option pricing. I am grateful to Alex d'Aspremont for pointing out that the option pricing problem provides a nice application of duality theory. Finally, I'd like to acknowledge the fact that half (four out of eight) of the typos were reported to me by Trond Steihaug. Thanks Trond!

Princeton, NJ, USA

Robert J. Vanderbei

Preface to 4th Edition

Besides the ongoing tweaking and refining of the language and presentation of the material, this edition also features new material in Chapters 4 and 12 on the average performance of the simplex method.

I'd like to thank Cagin Ararat and Firdevs Ulus for carefully reviewing and commenting on this new material.

Princeton, NJ, USA

Robert J. Vanderbei

Contents

Preface	vii
Preface to 2nd Edition	xi
Preface to 3rd Edition	xiii
Preface to 4th Edition	xv
Part 1. Basic Theory: The Simplex Method and Duality	1
Chapter 1. Introduction	3
1. Managing a Production Facility	3
2. The Linear Programming Problem	5
Exercises	7
Notes	9
Chapter 2. The Simplex Method	11
1. An Example	11
2. The Simplex Method	14
3. Initialization	16
4. Unboundedness	18
5. Geometry	19
Exercises	20
Notes	23
Chapter 3. Degeneracy	25
1. Definition of Degeneracy	25
2. Two Examples of Degenerate Problems	25
3. The Perturbation/Lexicographic Method	28
4. Bland's Rule	31
5. Fundamental Theorem of Linear Programming	33
6. Geometry	33
Exercises	36
Notes	37
	xvii

Chapter 4. Efficiency of the Simplex Method	39
1. Performance Measures	39
2. Measuring the Size of a Problem	39
3. Measuring the Effort to Solve a Problem	40
4. Worst-Case Analysis of the Simplex Method	41
5. Empirical Average Performance of the Simplex Method	44
Exercises	50
Notes	52
Chapter 5. Duality Theory	53
1. Motivation: Finding Upper Bounds	53
2. The Dual Problem	54
3. The Weak Duality Theorem	55
4. The Strong Duality Theorem	57
5. Complementary Slackness	63
6. The Dual Simplex Method	64
7. A Dual-Based Phase I Algorithm	66
8. The Dual of a Problem in General Form	67
9. Resource Allocation Problems	69
10. Lagrangian Duality	72
Exercises	73
Notes	79
Chapter 6. The Simplex Method in Matrix Notation	81
1. Matrix Notation	81
2. The Primal Simplex Method	83
3. An Example	86
4. The Dual Simplex Method	90
5. Two-Phase Methods	92
6. Negative Transpose Property	93
Exercises	95
Notes	97
Chapter 7. Sensitivity and Parametric Analyses	99
1. Sensitivity Analysis	99
2. Parametric Analysis and the Homotopy Method	102
3. The Parametric Self-Dual Simplex Method	105
Exercises	107
Notes	109
Chapter 8. Implementation Issues	111
1. Solving Systems of Equations: LU -Factorization	111
2. Exploiting Sparsity	115
3. Reusing a Factorization	119
4. Performance Tradeoffs	123
5. Updating a Factorization	124

6. Shrinking the Bump	127
7. Partial Pricing	128
8. Steepest Edge	129
Exercises	131
Notes	132
Chapter 9. Problems in General Form	133
1. The Primal Simplex Method	133
2. The Dual Simplex Method	135
Exercises	140
Notes	140
Chapter 10. Convex Analysis	141
1. Convex Sets	141
2. Carathéodory's Theorem	143
3. The Separation Theorem	144
4. Farkas' Lemma	146
5. Strict Complementarity	147
Exercises	149
Notes	150
Chapter 11. Game Theory	151
1. Matrix Games	151
2. Optimal Strategies	153
3. The Minimax Theorem	155
4. Poker	157
Exercises	161
Notes	163
Chapter 12. Regression	165
1. Measures of Mediocrity	165
2. Multidimensional Measures: Regression Analysis	167
3. L^2 -Regression	168
4. L^1 -Regression	170
5. Iteratively Reweighted Least Squares	171
6. An Example: How Fast Is the Simplex Method?	173
Exercises	178
Notes	183
Chapter 13. Financial Applications	185
1. Portfolio Selection	185
2. Option Pricing	190
Exercises	194
Notes	195

Part 2. Network-Type Problems	197
Chapter 14. Network Flow Problems	199
1. Networks	199
2. Spanning Trees and Bases	202
3. The Primal Network Simplex Method	206
4. The Dual Network Simplex Method	211
5. Putting It All Together	214
6. The Integrality Theorem	215
Exercises	216
Notes	224
Chapter 15. Applications	225
1. The Transportation Problem	225
2. The Assignment Problem	227
3. The Shortest-Path Problem	228
4. Upper-Bounded Network Flow Problems	231
5. The Maximum-Flow Problem	233
Exercises	235
Notes	239
Chapter 16. Structural Optimization	241
1. An Example	241
2. Incidence Matrices	243
3. Stability	244
4. Conservation Laws	245
5. Minimum-Weight Structural Design	249
6. Anchors Away	250
Exercises	252
Notes	254
Part 3. Interior-Point Methods	255
Chapter 17. The Central Path	257
Warning: Nonstandard Notation Ahead	257
1. The Barrier Problem	257
2. Lagrange Multipliers	259
3. Lagrange Multipliers Applied to the Barrier Problem	262
4. Second-Order Information	264
5. Existence	264
Exercises	266
Notes	267

Chapter 18. A Path-Following Method	269
1. Computing Step Directions	269
2. Newton's Method	270
3. Estimating an Appropriate Value for the Barrier Parameter	272
4. Choosing the Step Length Parameter	272
5. Convergence Analysis	274
Exercises	279
Notes	283
Chapter 19. The KKT System	285
1. The Reduced KKT System	285
2. The Normal Equations	286
3. Step Direction Decomposition	288
Exercises	290
Notes	291
Chapter 20. Implementation Issues for Interior-Point Methods	293
1. Factoring Positive Definite Matrices	293
2. Quasidefinite Matrices	296
3. Problems in General Form	302
Exercises	307
Notes	307
Chapter 21. The Affine-Scaling Method	309
1. The Steepest Ascent Direction	309
2. The Projected Gradient Direction	311
3. The Projected Gradient Direction with Scaling	312
4. Convergence	316
5. Feasibility Direction	317
6. Problems in Standard Form	319
Exercises	320
Notes	321
Chapter 22. The Homogeneous Self-Dual Method	323
1. From Standard Form to Self-Dual Form	323
2. Homogeneous Self-Dual Problems	324
3. Back to Standard Form	334
4. Simplex Method vs. Interior-Point Methods	336
Exercises	339
Notes	341
Part 4. Extensions	343
Chapter 23. Integer Programming	345
1. Scheduling Problems	345

2. The Traveling Salesman Problem	346
3. Fixed Costs	349
4. Nonlinear Objective Functions	350
5. Branch-and-Bound	351
Exercises	361
Notes	362
Chapter 24. Quadratic Programming	363
1. The Markowitz Model	363
2. The Dual	367
3. Convexity and Complexity	370
4. Solution via Interior-Point Methods	373
5. Practical Considerations	374
Exercises	376
Notes	378
Chapter 25. Convex Programming	379
1. Differentiable Functions and Taylor Approximations	379
2. Convex and Concave Functions	380
3. Problem Formulation	380
4. Solution via Interior-Point Methods	381
5. Successive Quadratic Approximations	382
6. Merit Functions	383
7. Parting Words	385
Exercises	385
Notes	388
Erratum	E1
Appendix A. Source Listings	389
1. The Self-Dual Simplex Method	390
2. The Homogeneous Self-Dual Method	393
Answers to Selected Exercises	395
Bibliography	399
Index	407

The original version of this book was revised.

An erratum to this book can be found at DOI [10.1007/978-1-4614-7630-6_26](https://doi.org/10.1007/978-1-4614-7630-6_26)