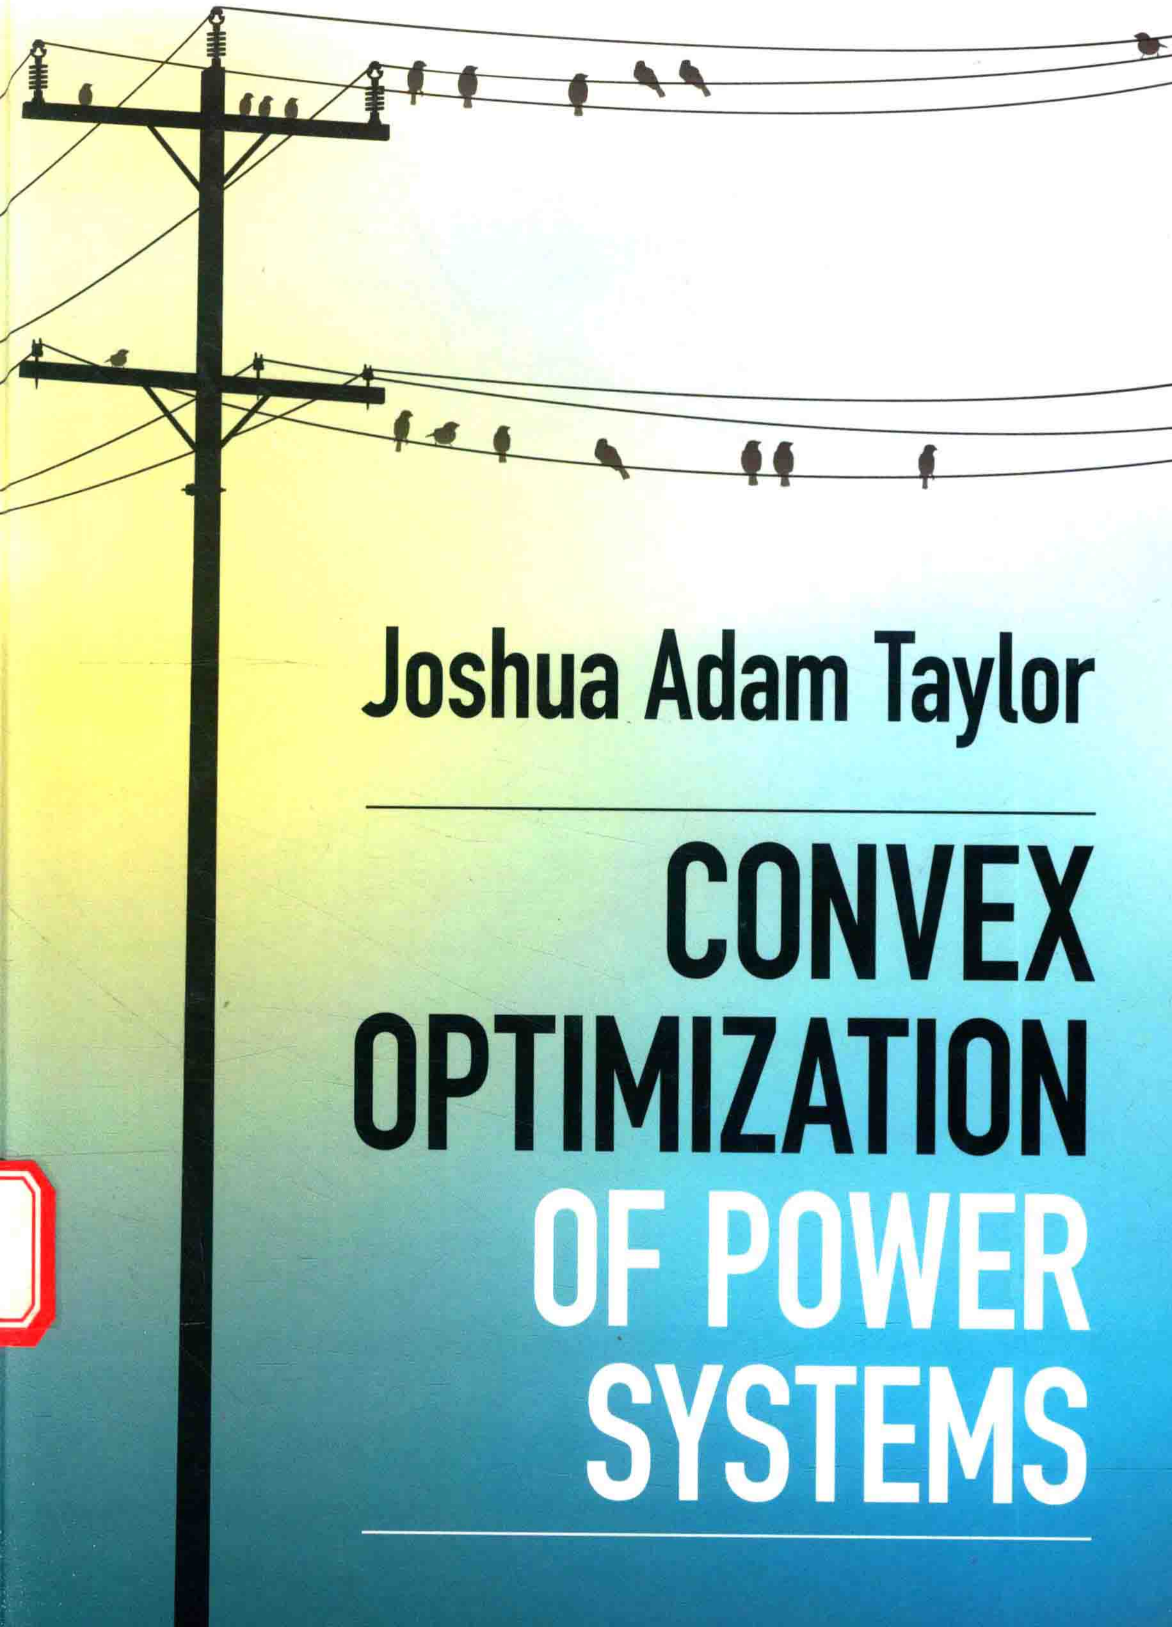

Joshua Adam Taylor

CONVEX OPTIMIZATION OF POWER SYSTEMS

Convex Optimization of Power Systems

JOSHUA ADAM TAYLOR
University of Toronto

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781107076877

First published 2015

Printed in the United Kingdom by TJ International Ltd, Padstow, Cornwall

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

Library of Congress Cataloguing in Publication data
Taylor, Joshua Adam, 1983–
Convex optimization of power systems / Joshua Adam Taylor.
pages cm
ISBN 978-1-107-07687-7 (hardback)
1. Electric power systems – Mathematical models. 2. Electric power distribution – Mathematics. 3. Convex programming. 4. Mathematical optimization. I. Title.
TK1005.T427 2015
621.3101′5196–dc23

2014032238

ISBN 978-1-107-07687-7 Hardback

Convex Optimization of Power Systems

Optimization is ubiquitous in power system engineering. Drawing on powerful, modern tools from convex optimization, this rigorous exposition introduces essential techniques for formulating linear, second-order cone, and semidefinite programming approximations to the canonical optimal power flow problem, which lies at the heart of many different power system optimizations.

Convex models in each optimization class are then developed in parallel for a variety of practical applications such as unit commitment, generation and transmission planning, and nodal pricing. Presenting classical approximations and modern convex relaxations side-by-side, and a selection of problems and worked examples, this book is an invaluable resource for students and researchers from industry and academia in power systems, optimization, and control.

Joshua Adam Taylor is Assistant Professor of Electrical and Computer Engineering at the University of Toronto.

For Elodie

Preface

The application of optimization to power systems has become so common that it deserves treatment as a distinct subject. The abundance of optimization problems in power systems can give the impression of diversity, but in truth most are merely layers on a common core: the steady-state description of power flow in a network. In this book, many of the most prominent examples of optimization in power systems are unified under this perspective.

As suggested by the title, this book focuses exclusively on convex frameworks, which by reputation are phenomenally powerful but often too restrictive for realistic, nonconvex power system models. In Chapter 3, the application of classical and recent mathematical techniques yields a rich spectrum of convex power flow approximations ranging from high tractability and low accuracy to slightly reduced tractability and high accuracy. The remaining chapters explore problems in power system operation, planning, and economics, each consisting of details layered on top of the convex power flow approximations. Because all formulations can be solved using standard software packages, only models are presented, which is a departure from most books on power systems. It is a major perk of convex optimization that the user often does not need to program an algorithm to proceed.

I should comment that this book is not an up-to-date exposition of power system applications or optimization theory and that, inevitably, many important topics in both fields have been omitted. My intention has rather been to bridge modern convex optimization and power systems in a rigorous manner. While I have attempted to be mathematically self-contained, the pace assumes an advanced undergraduate level of mathematical exposure (linear algebra, calculus, and some probability) as well as familiarity with power systems and optimization. This book could be used in a course on power system optimization or as a mathematical supplement to a course in power system design, operation, or economics. It is my hope that it will also prove useful to researchers in power systems with an interest in optimization and vice versa, and to industry practitioners seeking firm foundations for their optimization applications.

Acknowledgments

I started this book in the fall of 2011. Most of the present content I learned under the guidance of Franz Hover during my graduate school years at the Massachusetts Institute of Technology and of Duncan Callaway, Kameshwar Poolla, and Pravin Variaya during my postdoctoral studies at the University of California, Berkeley. Certainly, without the open-minded environments they created, this project would have never been attempted. I must also thank a number of colleagues whom I've benefitted from regular discussions with: Eilyan Bitar, Brendan Englot, Reza Iravani, Deepa Kundur, Johanna Mathieu, Daniel Muenz, Ashutosh Nayyar, Andy Packard, Matias Negrete-Pincetic, Anand Subramanian, and many others who have made power systems, control, and optimization such pleasant fields to work in. Finally, Julie Lancashire and Sarah Marsh at Cambridge University Press have made the final stages of this book a highly enjoyable process.

Notation

AC	Alternating current
DC	Direct current
LP	Linear programming
QP	Quadratic programming
SOC(P)	Second-order cone (programming)
SD(P)	Semidefinite (programming)
(C)QCP	(Convex) quadratically constrained programming
MI	Mixed integer
NLP	Nonlinear programming
KKT	Karush-Kuhn-Tucker (conditions)
PNE	Pure strategy Nash equilibrium
MNE	Mixed strategy Nash equilibrium

i	$\sqrt{-1}$
$\mathbb{R}$	The set of real numbers
$\mathbb{C}$	The set of complex numbers
$\mathbb{Z}$	The set of integers
x_i	The i^{th} entry of the vector x
x^k	The k^{th} version of the quantity x, typically corresponding to the k^{th} scenario or time period
$\mathrm{Re}\, x$	The real part of x
$\mathrm{Im}\, x$	The imaginary part of x
$\|x\|$	The absolute value of x
$\|\|x\|\|$	The two-norm of x, $\sqrt{\sum_i x_i^2}$
X_{ij}	The entry at the i^{th} row and j^{th} column of the matrix X
X^T	The transpose of X
X^*	The Hermitian transpose of X. When X is scalar, the complex conjugate.
$X \succeq 0$	The matrix X is positive semidefinite.
$\mathrm{rank}\, X$	The rank of X
$\mathrm{tr}\, X$	The trace of X, $\sum_i X_{ii}$
$\det X$	The determinant of X
∇	The gradient operator

To condense exposition, this book employs somewhat relaxed indexing notation. Because there is little risk of ambiguity, i will often be used simultaneously as the imaginary unit and as an index, for example iq_i would be $\sqrt{-1}$ times the i^{th} entry of q. In most cases, constraint indexing will not be explicitly declared; for example,

$$g_i(x) \leq 0$$

is implicitly enforced over $i = 1, \ldots, n$, which is almost always the set of nodes in the network. Similarly, the sum

$$\sum_{ij} x_{ij}$$

is over all relevant node pairs ij, which are usually those connected by lines. Indexing is denoted explicitly when it is not over a standard set, such as when summing over a subset of nodes.

This book makes extensive use of *feasible sets* as organizational tools. Given a collection of constraints $g_i(x) \leq 0$, the corresponding feasible set is

$$\{x \mid g_i(x) \leq 0\},$$

i.e., the set of points for which every constraint is satisfied.

Contents

1 Introduction

1.1 Recent history

Streetlights, subways, the Internet, this book you are reading now – it is difficult to imagine life without such amenities, all enabled by electric power. To support such a vast set of technologies, electric power systems have grown into some of the most complex and expensive machines in existence. While much of this growth resembles an organic process more than deliberate design, the advent of computing is enabling us more and more to direct the evolution of power systems toward greater efficiency, reliability, and versatility.

At the time of writing, the complexity of power systems is poised to take off. This is largely due to shifts toward renewable energy production and the active involvement of power consumers through demand response, as well as our still-developing handle on economic deregulation. To meet these challenges, new computational tools will be developed, and the most ubiquitous computation in power systems is optimization. An objective of this book is to simplify and unify various topics in power system optimization so as to provide a firm foundation for future developments.

At the heart of most power system optimizations are the equations of the steady-state, single-phase approximation to alternating current power flow in a network. Well-known problems like optimal power flow, reconfiguration, and transmission planning all consist of details layered on top of power flow. Nodal prices, a core component of electricity markets, are obtained from the dual of optimal power flow. It is therefore most unfortunate that the power flow equations are nonconvex, making all of these optimizations extremely difficult. We are thus faced with a tradeoff between realistic models that are too hard to solve at practical scales and tractable approximations.

For many years, linear programming (LP) was the most general efficiently solvable optimization class, and so many large-scale power system models were based on linear power flow approximations or even simpler descriptions like network flow or a real power balance. At the other extreme, a number of nonlinear programming (NLP) algorithms were developed for exact, nonconvex models. These approaches invariably encountered difficulty scaling to larger problem sizes due to the underlying NP-hardness of nonconvex optimization. This led some to resort to so-called metaheuristic algorithms, which make little use of problem structure and give little indication of their performance. Beyond their scalability issues, NLP and metaheuristic approaches can be tiresome to implement because they often require the user to program both the mathematical model and algorithm. On the other hand, one would rarely write their own algorithm to solve an LP because of the many available professional-grade commercial and academic implementations.

In 1984, there was a turning point for convex optimization when Karmarkar invented the first practical, polynomial-time interior point method for LP [1]. Over the next decade, second-order cone programming (SOCP) and semidefinite programming (SDP) emerged as convex generalizations of LP also admitting polynomial-time interior point methods [2]. Of equal importance, SOCP and SDP are now featured in a number of standard software packages, making them similarly user friendly.

The enhanced modeling capabilities of SOCP and SDP brought about an explosion of research applications, some of which can be found in the standard text [3]. In 2006, ripples from the previous twenty years were felt in power systems when power flow in radial networks was posed as an SOCP in Jabr [4] and again in 2008 when an SDP approximation of optimal power flow was developed in Bai, Wei, Fujisawa, and Wang [5]. A substantial body of research has materialized in the short time since then, both theoretically characterizing the new SOCP and SDP power flow approximations and applying them in a variety of power system contexts. For most power system optimization problems, we now have a spectrum of LP, SOCP, SDP, and NLP models to choose from, each with a different balance of realism and scalability.

1.2 Structure and outline

This book only contains models and, with the exception of Chapters 2 and 7, rarely mentions algorithms. This is *not* because the algorithms are not worth knowing or decoupled from modeling; one can *always* do better by formulating optimization models and algorithms jointly. Rather, here this wisdom is applied by formulating models so that they can be solved by certain algorithms. This approach is a luxury we can afford because optimization is a relatively mature field: for a desired level of scalability, it identifies the corresponding tradeoff between efficiency and descriptiveness. Here, this manifests as a hierarchy of convex optimization classes, the main elements of which are LP, SOCP, and SDP. LP is the most efficient and least descriptive, SDP vice versa, and SOCP is in between. As discussed in the previous section, once a model has been formulated within one of these classes, it can be conveniently solved using standard software packages. By tailoring our models to these classes, we arrive at tractable formulations far more easily than if we were to design both model and algorithm from scratch.

The resulting separation between models and algorithms should not be seen as a restriction but as a starting point for further specialization and extension. For example, once a problem has been formulated as an LP, uncertainty can be mechanistically incorporated using robust optimization (Section 7.1.2), or the problem can be split into different chunks for multiple processors or agents (Section 7.2).

A central motif in this book is posing complicated models as layers on top of more basic ones. To avoid rewriting the same constraints over and over, this book makes extensive use of *feasible sets* to package frequently occurring groups of constraints for concise representation in other problems. Consequently, certain parts of this book are highly cumulative. Each chapter is concluded by a summary highlighting important points and open problems, as well as a small selection of exercises. In addition to those given, which are all analytical, students are of course encouraged to implement each chapter's models and examples using their preferred optimization platform.

The chapter structure is summarized below.

Chapter 2: This chapter provides a minimal introduction to optimization and power system modeling. In particular, it defines LP, SOCP, and SDP, and the tools used to construct convex relaxations. It also derives the quadratic steady-state, single-phase approximation to power flow in a network.

Chapter 3: This chapter defines the basic, nonconvex optimal power flow problem. It then derives classical linear approximations and modern SOC and SD relaxations. The constraints in these models form the foundation of all subsequent chapters in this book.

Chapter 4: This chapter constructs linear, SOC, and SD versions of a number of central optimization problems in power system operations using the approximations from the previous chapter. Here and in the next chapter we encounter a number of mixed-integer constraints, which make these problems challenging even with linear power flow constraints. This chapter also takes brief detours though inventory control and linear quadratic regulation.

Chapter 5: Similar to the last chapter, this chapter constructs infrastructure planning problems around the power flow approximations of Chapter 3. In this chapter, every problem is extremely difficult due to the integer constraints required to describe component installation.

Chapter 6: This chapter discusses electricity markets, in which *nodal prices* are obtained from the dual of optimal power flow. Because each convex approximation has strong duality, prices are guaranteed to support *economic dispatch*. This chapter also discusses why basic economic assumptions never hold in practice and briefly summarizes some game theoretic analyses of market power.

Chapter 7: This final chapter surveys some promising directions for future work.

1.3 On approximations

Every model in this book is an approximation. In fact, every mathematical model ever is an approximation, which inevitably fails to capture some fine physical detail. However, it is worth stating this explicitly here because every model in this book involves an approximation of one particular model: the steady-state, single-phase description of electric power in a network of conductors. We derive this model in Section 2.5 and place it in the context of optimization in Section 3.1.

The steady-state, single-phase description of electric power is a special approximation because it derives from very natural physical assumptions and enables power system engineers to take a large step from simulation to design, which is what this book is all about. By definition, optimization is the highest (mathematical) form of design. Unfortunately, the steady-state description of electric power isn't quite right for optimization because it is nonconvex. This book attempts to make the mildest further adjustments necessary for this model and those built atop it to enjoy all that optimization has to offer.

With this perspective in mind, it is helpful to remember the following two statements when using this book.

- We would always use a more realistic description of electric power like an unbalanced steady-state or transient model were it practical to do so. So, when it is practical, use one of them and not the convex approximations in this book. We briefly elaborate on this in Section 3.1.1.
- While immensely important, the steady-state description of electric power is not sacrosanct. It is an approximation like all of the other models in this book, which just happens to be (slightly) closer to reality. From this perspective, the convex approximations in this book are no less valid than the model from which they are derived.

References

[1] N. Karmarkar, "A new polynomial-time algorithm for linear programming," in *Proceedings of the Sixteenth Annual ACM Symposium on Theory of Computing*, ser. STOC '84. New York: ACM, 1984, pp. 302–311.

[2] Y. Nesterov and A. Nemirovski, "Interior point polynomial methods in convex programming," *SIAM Studies in Applied Mathematics*, vol. 13, 1994.

[3] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York: Cambridge University Press, 2004.

[4] R. Jabr, "Radial distribution load flow using conic programming," *IEEE Transactions on Power Systems*, vol. 21, no. 3, pp. 1458 –1459, Aug. 2006.

[5] X. Bai, H. Wei, K. Fujisawa, and Y. Wang, "Semidefinite programming for optimal power flow problems," *International Journal of Electrical Power and Energy Systems*, vol. 30, no. 6–7, pp. 383–392, 2008.

2 Background

This chapter summarizes the basic technical concepts used throughout this book. As stated in the Introduction, this book focuses on modeling, so most algorithmic aspects are left "under the hood." Because this book is intended to appeal to anyone familiar with power systems or optimization, background material on both topics is covered, albeit at the minimum depth necessary to access the later material.

2.1 Convexity and computational complexity

We begin with a few core concepts. A point $x_0 \in X$ is a *global minimum* of the function $f(x)$ over the set $X \subseteq \mathbb{R}^n$ if $f(x_0) \leq f(x)$ for all $x \in X$. If $f(x)$ is continuous and X is compact, which is to say closed and bounded, such a point is guaranteed to exist. x_0 is a *local minimum* of $f(x)$ if there exists an $\epsilon > 0$ for which $f(x_0) \leq f(x)$ for all $x \in X$ satisfying $\|x - x_0\| \leq \epsilon$. All minima of a convex function achieve the same function value and are therefore global. In general, a function may have multiple local and global minima.

To find a global minimum of a convex function, choose a descending algorithm, let it run free, and in a perfect world it will eventually end up there. (In the real world, large problem sizes or bad numerical conditioning can derail any algorithm.) The intuitive simplicity of convexity translates to a genuine computational advantage, which is evinced by the powerful algorithms that exist for convex optimization and the extreme difficulty of nonconvex optimization. This section gives some basic characterizations of convexity and describes the varieties of convex optimization problems encountered in this book. For more comprehensive coverage, the reader is referred to endnotes [1–5], and to endnotes [6, 7] for theoretical treatments of convex functions and sets.

A function f is convex if, for any two points in its domain, x and y,

$$f(\alpha x + (1-\alpha)y) \geq \alpha f(x) + (1-\alpha)f(y) \quad \text{for all } \alpha \in [0,1].$$

This means that any point on the straight line between $(x, f(x))$ and $(y, f(y))$ is greater than or equal to the function value at the corresponding point between x and y. When f is twice-differentiable, it is convex if and only if its Hessian is positive semidefinite:

$$\nabla^2 f(x) \succeq 0.$$