

Lane A. Hemaspaandra
Mitsunori Ogihara

The Complexity Theory Companion



Texts in Theoretical Computer Science

An EATCS Series

Editors: W. Brauer G. Rozenberg A. Salomaa
On behalf of the European Association
for Theoretical Computer Science (EATCS)

Advisory Board: G. Ausiello M. Broy S. Even
J. Hartmanis N. Jones T. Leighton M. Nivat
C. Papadimitriou D. Scott

Springer-Verlag Berlin Heidelberg GmbH

Lane A. Hemaspaandra • Mitsunori Ogihara

The Complexity Theory Companion

With 43 Figures



Springer

Authors

Prof. Dr. Lane A. Hemaspaandra
Prof. Dr. Mitsunori Ogihara
Department of Computer Science
Rochester, NY 14627
USA
{lane,ogihara}@cs.rochester.edu

Series Editors

Prof. Dr. Wilfried Brauer
Institut für Informatik
Technische Universität München
Arcisstrasse 21, 80333 München, Germany
brauer@informatik.tu-muenchen.de

Prof. Dr. Grzegorz Rozenberg
Leiden Institute of Advanced Computer Science
University of Leiden
Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
rozenber@liacs.nl

Prof. Dr. Arto Salomaa
Data City
Turku Centre for Computer Science
20500 Turku, Finland
asalomaa@utu.fi

Library of Congress Cataloging-in-Publication Data applied for

Die Deutsche Bibliothek CIP-Einheitsaufnahme

Hemaspaandra, Lane A.:
The complexity theory companion/Lane A. Hemaspaandra; Mitsunori Ogihara.
Berlin; Heidelberg; New York; Barcelona; Hong Kong; London; Milan;
Paris; Tokyo: Springer, 2002
(Texts in theoretical computer science)

ACM Computing Classification (1998): F.1, I.2.2, I.4

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 4, 1955, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

ISBN 978-3-642-08581-7 ESDN 978-3-662-04330-1 (eBook)
DOI 10.1007/978-3-662-04330-1

© Springer-Verlag Berlin Heidelberg 2002

Originally published by Springer-Verlag Berlin Heidelberg New York in 2002.

Softcover reprint of the hardcover 1st edition 2002

The use of general descriptive names, trademarks, 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.

Cover Design: Künzel/Opka, Heidelberg

Typesetting: Camera ready by the authors

Printed on acid-free paper SPIN 10/20529 450142SR - 5 4 3 2 1 0

This book is dedicated to our families—the best of companions.

Preface

Invitation

Secret 1 *Algorithms are at the heart of complexity theory.*

That is the dark secret of complexity theory. It is recognized by complexity theorists, but would be literally incredible to most others. In this book, we hope to make this secret credible. In fact, the real secret is even more dramatic.

Secret 2 *Simple algorithms are at the heart of complexity theory.*

A corollary of Secret 2 is that every practitioner of computer science or student of computer science already possesses the ability required to understand, enjoy, and employ complexity theory.

We realize that these secrets fly in the face of conventional wisdom. Most people view complexity theory as an arcane realm populated by pointy-hatted (if not indeed, pointy-headed) sorcerers stirring cauldrons of recursion theory with wands of combinatorics, while chanting incantations involving complexity classes whose very names contain hundreds of characters and bear the tongue of mere mortals. This stereotype has sprung up in part due to the small amount of esoteric research that fits this bill, but the stereotype is more strongly attributable to the failure of complexity theorists to communicate in expository forums the central role that algorithms play in complexity theory.

Throughout this book—from the tree-pruning and interval-pruning algorithms that shape the first chapter to the query simulation procedures that dominate the last chapter—we will see that proofs in complexity theory usually employ algorithms as their central tools. In fact, to more clearly highlight the role of algorithmic techniques in complexity theory, *this book is organized by technique rather than by topic*. That is, in contrast to the organization of other books on complexity theory, each chapter of this book focuses on one technique—what it is, and what results and applications it has yielded.

The most thrilling times in complexity theory are when a new technique is introduced and sweeps like fire over the field. In addition to highlighting the centrality of algorithms in the proof arsenal of complexity theory, we feel that our technique-based approach more vividly conveys to the reader the flavor and excitement of such configurations. We invite the reader to come

with us as we present nine techniques, usually simple and algorithmic, that burned away some of the field's ignorance and helped form the landscape of modern complexity theory.

Usage

We intend this book as a companion for students and professionals who seek an accessible, algorithmically oriented, research-oriented, up-to-date guide to some of the most interesting techniques of complexity theory. The authors and their colleague Joel Seiferas have test-driven the book's approach in two different courses at the University of Rochester. We have used this techniques-based approach in Rochester's one-semester basic complexity theory course, which is taken by all first-year computer science graduate students and also by those undergraduates specializing or specially interested in theoretical computer science, and in our second course on complexity theory, which is taken by all second-year graduate students as their theory "breadth" course.

We found in both these course settings that the technique-based approach allowed us to impart to students a significant amount of the feel and experience of complexity theory research and led to more student interest and involvement than occurred in earlier course incarnations using other texts. We expect that this will not only benefit the complexity theory students in the courses, but will also help all the course's students become prepared to do work that is theoretically aware, informed, and well-grounded.

At times, we stop the flow of a proof or discussion with a "Pause to Ponder." These are places at which we encourage the reader to pause for a moment and find his or her own solution to the issue raised. Even an unsuccessful attempt to craft a solution will usually make the proof/discussion that follows clearer and more valuable, as the reader will better understand the challenges posed by the hurdle that the proof/discussion overcomes.

With some exceptions due to result dependencies, the non-appendix chapters are generally ordered to put the easier chapters near the start of the book and the more demanding chapters near the end of the book.

Acknowledgments

We are extraordinarily indebted to the following people, who proofread one or more chapters, for their invaluable help, suggestions, corrections, and insights: Eric Allender, Russell Bent, Alim Beygelzimer, Matthew Boutsell, Samuel Chen, Yin-Ho Chung, Louis Demet, Gregory Goldstein, Fred Green, Ulrich Hertrampf, Chris Homer, Gabriel Istrate, Jason Ku, David Lagakos, Andrew Learn, Tao Li, Ioan Macarie, Proshanto Mukherji, Kenneth Regan, William Scherer III, Alan Selman, D. Sivakumar, Howard Strackling, Robert Swier,

Meyur Thakur, Jonathan Turner, Jacopo Torán, Leon Torenvliet, Dieter van Melkebeek, Heribert Vollmer, Julie Zhong, and Marius Zimand. We also thank the many other people who have helped us with advice, comments, corrections, literature pointers, most-recent-version information, and suggestions: Andris Ambainis, Vikraman Arvind, Richard Beigel, Nate Blaylock, Daniel Boyer, Jin-Yi Cai, Diane Cass, Stephen Chong, Liana Furtlow, William Gasarch, Viliam Geffert, Oded Goldreich, Juris Hartmanis, Edith Hemaspaandra, Paul Hord, Sven Kosub, Richard Lipton, Alexis Maciel, Wolfgang Mecke, Christos Papadimitriou, Thanos Papatheodorou, Eduardo Pinheiro, Robert Rottlinger, Jörg Rothe, Alex Samorodnitsky, Marius Schaefer, Michael Scheer, Uwe Schöning, Joel Sciferas, Samik Sengupta, Carl Smith, Scott Stinson, Madhu Sudan, Chongjiao Tang, Jun Tian, Thomas Thierauf, Luca Trevisan, Chris Umans, Osamu Watanabe, Choo Yap, and Steffen Zilles. Any remaining errors are the responsibility of the authors.

We thank our thesis advisors, Juris Hartmanis and Kojiro Kobayashi; their insightful, dedicated, joyous approach to research has been a continuing inspiration to us.

We appreciate the grants—NSF-CCR-8957604, NSF-INT-9118781/ISPS ENGR-207, NSF-CCR-9322513, NSF-INT-9513308/DAAD-315-PRO-60-95, NSF-CCR-9701911, NSF-CCR-9725021, NSF-INT-9726724, NSF-INT-9815095/DAAD-315-PPP-gü-98, NSF-DUE-9980941, DARPA-F30602-98-2-0133, NIA-R01-AG-18231—that have supported our research programs during the planning and writing of this book.

For generous advice, help, and support, we thank the Springer-Verlag series editors and staff, namely, Wilfried Brauer, Geroport Rosenberg, Arto Salomaa, Alfred Hofmann, Frank Holte, Ingeborg Mayer, Sherryl Studdell, and Hans Wössner. Our own department's technical and secretarial staff—especially Jill Furster, Elaine Heberle, and Jim Roche—was invaluable in keeping our computers up and our manuscript copied and circulating, and we much appreciate their help.

We are grateful to those colleagues—Peter van Emde Boas, Harald Hempel, Jörg Rothe, Alan Selman, Shunsuke Toda, Leon Torenvliet, Heribert Vollmer, Klaus Wagner, Osamu Watanabe, and Gerald Wechsung—who have generously hosted our research visits during the planning and writing of this book, and to the many colleagues, as cited in the Bibliographic Notes sections, who have collaborated with us on the research described in some sections of this book.

Above all, we thank Elinh, Ellen, Emil, and Kira for their advice, love, and support.

Lars A. Hemaspaandra

Mitsuru Oshima

Rochester, NY

October 2001

Contents

Preface	vii
Introduction	vii
Usage	viii
1. The Self-Reducibility Technique	1
1.1 CEM: There Are No Sparse NP-Complete Sets Unless $P=NP$..	2
1.2 The Turing Case	18
1.3 The Case of Merely Putting Sparse Sets in NP = P: The Hartmanis Immerman Sewelson Encoding	22
1.4 OPEN ISSUE: Does the Disjunctive Case Hold?	26
1.5 Bibliographic Notes	26
2. The One-Way Function Technique	31
2.1 GEM: Characterizing the Existence of One-Way Functions ..	32
2.2 Unambiguous One-Way Functions Exist If and Only If Bounded-Ambiguity One-Way Functions Exist	35
2.3 Strong, Total, Commutative, Associative One-Way Functions Exist If and Only If One-Way Functions Exist	36
2.4 OPEN ISSUE: Low-Ambiguity, Commutative, Associative One-Way Functions?	42
2.5 Bibliographic Notes	43
3. The Tournament Divide and Conquer Technique	45
3.1 CEM: The Semi-Feasible Sets Have Small Circuits	45
3.2 Optimal Advice for the Semi-Feasible Sets	48
3.3 Unique Solutions Collapse the Polynomial Hierarchy	56
3.4 OPEN ISSUE: Are the Semi-Feasible Sets in P/log ?	63
3.5 Bibliographic Notes	63
4. The Isolation Technique	67
4.1 CEM: Isolating a Unique Solution	68
4.2 Toda's Theorem: $P \Pi \subseteq P^{PP}$	72
4.3 $NL/\text{poly} = UL/\text{poly}$	82

4.4	OPEN ISSUE: Do Ambiguous and Unambiguous Nondeterminism Coincide?	57
4.5	Bibliographic Notes	37
5.	The Witness Reduction Technique	91
5.1	Framing the Question: Is $\#P$ Closed Under Proper Subtraction?	91
5.2	GEM: A Complexity Theory for Feasible Closure Properties of $\#P$	93
5.3	Intermediate Potential Closure Properties	90
5.4	A Complexity Theory for Feasible Closure Properties of OptP	103
5.5	OPEN ISSUE: Characterizing Closure Under Proper Decrement	105
5.6	Bibliographic Notes	100
6.	The Polynomial Interpolation Technique	109
6.1	GEM: Interactive Protocols for the Permanent	110
6.2	Enumerators for the Permanent	119
6.3	$\text{IP} = \text{PSPACE}$	122
6.4	$\text{MIP} = \text{NEXP}$	133
6.5	OPEN ISSUE: The Power of the Prover	163
6.6	Bibliographic Notes	103
7.	The Nonsolvable Group Technique	167
7.1	GEM: Width 5 Branching Programs Capture Nonuniform NC^1	168
7.2	Width-5 Bottleneck Machines Capture PSPACE	176
7.3	Width-2 Bottleneck Computation	181
7.4	OPEN ISSUE: How Complex Is Majority-Based Probabilistic Symmetric Bottleneck Computation?	192
7.5	Bibliographic Notes	192
8.	The Random Restriction Technique	197
8.1	GEM: The Random Restriction Technique and a Polynomial-Size Lower Bound for Parity	197
8.2	An Exponential-Size Lower Bound for Parity	207
8.3	PH and PSPACE Differ with Probability One	215
8.4	Oracles That Make the Polynomial Hierarchy Infinite	232
8.5	OPEN ISSUE: Is the Polynomial Hierarchy Infinite with Probability One?	231
8.6	Bibliographic Notes	231

9. The Polynomial Technique	235
9.1 GIM: The Polynomial Technique.....	236
9.2 Closure Properties of PP.....	241
9.3 The Probabilistic Language Hierarchy Collapses.....	252
9.4 OPEN ISSUE: Is PP Closed Under Polynomial-Time Turing Reductions?.....	259
9.5 Bibliographic Notes.....	266
A. A Rogues' Gallery of Complexity Classes	263
A.1 P: Determinism.....	264
A.2 NP: Nondeterminism.....	266
A.3 Oracles and Relativized Worlds.....	268
A.4 The Polynomial Hierarchy and Polynomial Space: The Power of Quantifiers.....	270
A.5 E, NE, EXP, and NEXP.....	274
A.6 P/Tpoly: Small Circuits.....	276
A.7 L, NL, etc.: Logspace Classes.....	277
A.8 NC, AC, LOGCFL: Circuit Classes.....	279
A.9 UP, P=NP, and US: Ambiguity-Scandal Computation and Unique Computation.....	281
A.10 #P: Counting Solutions.....	286
A.11 ZPP, RP, coRP, and BPP: Error-Bounded Probabilism.....	288
A.12 PP, C=P, and SPP: Counting Classes.....	290
A.13 FP, NPSV, and NPMV: Deterministic and Nondeterministic Functions.....	291
A.14 P-Sub: Semi-Feasible Computation.....	294
A.15 #P, Mod#P: Modulo-Based Computation.....	297
A.16 SpanP, OptP: Output-Cardinality and Optimization Function Classes.....	297
A.17 IP and MIP: Interactive Proof Classes.....	299
A.18 PBP, SP, SSP: Branching Programs and Bottleneck Computation.....	300
B. A Rogues' Gallery of Reductions	305
B.1 Reduction Definitions: \leq_p^A , \leq_p^B	305
B.2 Shorthands: R and E.....	307
B.3 Facts about Reductions.....	307
B.4 Circuit-Based Reductions: NC^A and AC^A	308
B.5 Bibliographic Notes.....	308
References	309
Index	335

1. The Self-Reducibility Technique

A set S is *sparse* if it contains at most polynomially many elements at each length, i.e.,

$$|\{x \mid x \in S \wedge |x| = n\}| \leq p(n), \quad (1.1)$$

This chapter studies one of the oldest questions in computational complexity theory: Can sparse sets be NP-complete?

As we noted in the Preface, the proofs of most results in complexity theory rely on algorithms, and the proofs in this chapter certainly support that claim. In Sect. 1.1, we¹ will use a sequence of increasingly elaborate deterministic tree-pruning and interval-pruning procedures to show that sparse sets cannot be $\leq_{\text{m}}^{\text{P}}$ -complete, or even $\leq_{\text{m}}^{\text{P}}$ -hard, for NP unless $\text{P} = \text{NP}$. (The appendices contain definitions of and introductions to the reduction types, such as $\leq_{\text{m}}^{\text{P}}$ and $\leq_{\text{m}}^{\text{NP}}$, and the complexity classes, such as P and NP, that are used in this book.)

Section 1.2 studies whether NP can have $\leq_{\text{m}}^{\text{P}}$ -complete or $\leq_{\text{m}}^{\text{P}}$ -hard sparse sets. $\text{P}^{\text{NP}(\mathcal{O}(\log n))}$ denotes the class of languages that can be accepted by some deterministic polynomial-time Turing machine allowed at most $\mathcal{O}(\log n)$ queries to some NP oracle. In Sect. 1.2, we will—via binary search, self-reducibility algorithms, and nondeterministic algorithms—prove that sparse sets cannot be $\leq_{\text{m}}^{\text{P}}$ -complete for NP unless the polynomial hierarchy collapses to $\text{P}^{\text{NP}(\mathcal{O}(\log n))}$, and that sparse sets cannot be $\leq_{\text{m}}^{\text{P}}$ -hard for NP unless the polynomial hierarchy collapses to NP^{NP} .

As is often the case in complexity-theoretic proofs, we will typically use in the construction of our algorithms the hypothesis of the theorem that the algorithm is establishing (e.g., we will build a P algorithm for SAT, and will use in the algorithm the—literally hypothetical—sparse $\leq_{\text{m}}^{\text{P}}$ -complete set for NP). In fact, this “theorems via algorithms under hypotheses” approach is employed in each section of this chapter.

Furthermore, most of Sects. 1.1 and 1.2 are unified by the spirit of their algorithmic attack, which is to exploit the “(disjunctive) self-reducibility” of SAT—basically, the fact that a boolean formula is satisfiable \Leftrightarrow and only if either it is satisfiable with its first variable set to $\overline{\text{false}}$ or it is satisfiable with

¹ In this book, “we” usually refers to the authors and the readers as we travel together in our exploration of complexity theory.

its first variable set to True. A partial exception to the use of this attack in those sections is the left set technique, which we use in Sect. 1.1.2. This technique, while in some sense a veiled tree-pruning procedure inspired by a long line of self-reducibility-based tree-pruning procedures, adds a new twist to this type of argument, rather than being a direct invocation of SAT's self-reducibility.

Section 1.3 studies not whether there are sparse NP-complete sets, but rather whether $NP = P$ contains any sparse sets at all. Like the previous sections, this section employs explicit algorithmic constructions that themselves use objects hypothesized to exist by the hypotheses of the theorems for which they are providing proofs. The actual result we arrive at is that $NP = P$ contains sparse sets if and only if deterministic and nondeterministic exponential time differ.

Throughout this book, we will leave the type of quantified variables implicit when it is clear from context what that type is. For example, in equation 1.1, the " $(\forall n)$ " is implicitly " $(\forall n \in \{0, 1, 2, \dots\})$," and " $(\exists x)$ " is typically a shorthand for " $(\exists x \in \Sigma^*)$." We will use a colon to denote a constraint on a variable, i.e., " $(\forall x: R(x)) [S(x)]$ " means " $(\forall x) [R(x) \rightarrow S(x)]$," and " $(\exists x: R(x)) [S(x)]$ " means " $(\exists x) [R(x) \wedge S(x)]$." For any set A and any natural number n , we will use $A^{\leq n}$ to denote the strings of A that are of length at most n , and we will use A^n to denote the strings of A that are of length exactly n . Given a Turing machine M , we will use $L(M)$ to denote the language accepted by the machine (with respect to whatever the acceptance mechanism of the machine is).

1.1 GEM: There Are No Sparse NP-Complete Sets Unless $P = NP$

1.1.1 Setting the Stage: The Pruning Technique

Before we turn to Mahaney's Theorem— NP has sparse complete sets only if $P = NP$ —and its generalization to bounded-truth-table reductions, we first prove two weaker results that display the self-reducibility-based tree-pruning approach in a simpler setting. (Below, in a small abuse of notation we are taking " 1 " in certain places—such as in expressions like " 1^* "—as a shorthand for the regular expression representing the set $\{1\}$.)

Definition 1.1 A set T is a tally set exactly if $T \subseteq 1^*$.

Theorem 1.2 If there is a tally set that is \leq_m^p -hard for NP , then $P = NP$.

Corollary 1.3 If there is a tally set that is NP -complete, then $P = NP$.

We note in passing that if $P = NP$, then the singleton set $\{1\}$ is trivially both NP complete and $coNP$ complete. Thus, all the "if...then..." theorems

of this section (Sect. 1.1) have true converses. We state them in “if...then...” form to stress the interesting direction.

Proof of Theorem 1.2 Let T be a tally set that is \leq_m^p hard for NP. Since the NP-complete set

$$\text{SAT} = \{f \mid f \text{ is a satisfiable boolean formula}\}$$

is in NP and T is \leq_m^p hard for NP, it follows that $\text{SAT} \leq_m^p T$. Let g be a deterministic polynomial-time function many-one reducing SAT to T . Let k be an integer such that $(\forall x)(|g(x)| \leq |x|^k + k)$; since g is computable by some deterministic polynomial-time Turing machine, such a k indeed must exist since that machine outputs at most one character per step.

We now give, under the hypothesis of the theorem, a deterministic polynomial-time algorithm for SAT, via a simple tree-pruning procedure. The input to our algorithm is a boolean formula F . Without loss of generality, let its variables be v_1, \dots, v_m and let $m \geq 7$. We will denote the result of assigning values to some of the variables of F via expressions of the following form: $F[v_1 = \text{True}, v_3 = \text{False}]$, where True denotes the constant true and False denotes the constant false. For example, if $F = v_1 \vee v_2 \vee v_3$ then

$$F[v_1 = \text{True}, v_3 = \text{False}] = \text{True} \vee v_2 \vee \text{False},$$

and

$$(F[v_1 = \text{True}])(v_3 = \text{False}) = \text{True} \vee v_2 \vee \text{False}.$$

Our algorithm has stages numbered $0, 1, \dots, m-1$. At the end of each stage (except the final one), we pass forward a collection of boolean formulas. Initially, we view ourselves as having just completed Stage 0, and we view ourselves as passing forward from Stage 0 a collection, \mathcal{C} , containing the single formula F .

Stage i , $1 \leq i \leq m$, assuming that the collection at the end of Stage $i-1$ is the following collection of formulas: $\{F_1, \dots, F_\ell\}$.

Step 1 Let \mathcal{C} be the collection

$$\begin{aligned} & \{F_1[v_1 = \text{True}], F_2[v_1 = \text{True}], \dots, F_\ell[v_1 = \text{True}], \\ & F_1[v_1 = \text{False}], F_2[v_1 = \text{False}], \dots, F_\ell[v_1 = \text{False}]\}. \end{aligned}$$

Step 2 Set \mathcal{C}' to be \emptyset .

Step 3 For each formula f in \mathcal{C} (in arbitrary order) do:

1. Compute $g(f)$.
2. If $g(f) \in T$ and for no formula $h \in \mathcal{C}'$ does $g(f) = g(h)$, then add f to \mathcal{C}' .

End Stage i (\mathcal{C}' is the collection that gets passed on to Stage $i+1$)

The action of our algorithm at Stage $m-1$ is simple: F is satisfiable if and only if some member of the (variable-free) formula collection output by Stage m evaluates to being true.

As to the correctness of our algorithm, note that after Stage 0 it certainly holds that

$$\begin{aligned} \text{the collection, } C_0, \text{ contains some satisfiable formula} \\ \langle \rangle \\ F \text{ is satisfiable,} \end{aligned} \tag{1.2}$$

since after Stage 0 formula F is the only formula in the collection. Note also that, for each i , $1 \leq i \leq m$,

$$\begin{aligned} \text{The collection input to Stage } i \text{ contains some} \\ \text{satisfiable formula} \\ \langle \rangle \\ \text{The collection output by Stage } i \text{ contains some} \\ \text{satisfiable formula.} \end{aligned} \tag{1.3}$$

We now argue that this is so, via using a self-reducibility-based argument. In the present context, the relevant self-reducibility fact is that for any formula F containing v as one of its variables,

$$F \text{ is satisfiable} \iff ((F[v = \text{True}] \text{ is satisfiable}) \vee (F[v = \text{False}] \text{ is satisfiable})),$$

since any satisfying assignment must assign some value to each variable. So Step 1 of Stage i does no damage to our invariant, equation 1.3. What about Steps 2 and 3? (In terms of the connection to Step 1, it is important to keep in mind that if, for example, formula F having variable v is in our collection at the start of the stage and is satisfiable, then it must be the case that

$$(F[v = \text{True}] \text{ is satisfiable}) \vee (F[v = \text{False}] \text{ is satisfiable}),$$

so it must be the case that

$$g(F[v = \text{True}]) \in T \vee g(F[v = \text{False}]) \in T.$$

And of course, $T' \subset T$.) Steps 2 and 3 "prune" the formula set as follows. Each formula f from Step 1 is kept unless either

- a. $g(f) \notin T'$, or
- b. $g(f) \in T'$ but some $h \in C'$ has $g(f) = g(h)$.

Both these ways, (a) and (b), of dropping formulas are harmless. Recall that $\text{SAT} \leq_m^p T$ via function g , and so if $f \in \text{SAT}$ then $g(f) \in T$. However, regarding (a), $T' \subset T$ so if $g(f) \notin T'$ then $g(f) \notin T$, and so $f \notin \text{SAT}$. Regarding (b), if $g(f) = g(h)$ and h has already been added to the collection to be output by Stage (i), then there is no need to output f as—since $\text{SAT} \leq_m^p T$ via reduction g —we know that

$$f \in \text{SAT} \implies g(f) \in T$$

and

$$h \in \text{SAT} \iff g(h) \in T.$$

Thus, $f \in \text{SAT} \iff h \in \text{SAT}$, and so by discarding f but leaving in h , we do no damage to our invariant, equation 1.3. So by equations 1.2 and 1.3 we see that F is satisfiable if and only if some formula output by Stage m is satisfiable. As the formulas output by Stage m have no free variables, this just means that one of them must evaluate to being true, which is precisely what Stage $m + 1$ checks.

Thus, the algorithm correctly checks whether F is satisfiable. But is this a polynomial-time algorithm? $|F|$ will denote the length of F , i.e., the number of bits in the representation of F . Let $|F| = p$. Note that after any stage there are at most $p^k + k - 1$ formulas in the output collection, and each of these formulas is of size at most p . This size claim holds as each formula in an output collection is formed by one or more assignments of variables of F to being True or False, and such assignments certainly will not cause an increase in length. (In a standard, reasonable encoding). We will say that a string s is a *tally string* exactly if $s \in 1^*$. The $p^k + k - 1$ figure above holds as (due to the final part of Step 3 of our algorithm) we output at most one formula for each tally string to which ($n^k + k$ -time function) g can map, and even if g outputs a 1 on each step, g can output in $p^k + k$ steps no tally string longer than $1^{p^k + k}$. So, taking into account the fact that the empty string is a (degenerate) tally string, we have our $p^k + k + 1$ figure. From this, from the specification of the stages, and from the fact that g itself is a polynomial-time computable function, it follows clearly that the entire algorithm runs in time polynomial in $|F|$. \square

In the proof of Theorem 1.2, we used self-reducibility to split into two each member of a set of formulas, and then we pruned the resulting set using the fact that formulas mapping to non-tally strings could be eliminated, and the fact that only one formula mapping to a given tally string need be kept. By repeating this process we walked down the self-reducibility tree of any given formula, yet we pruned that tree well enough to ensure that only a polynomial number of nodes had to be examined at each level of the tree. By the self-reducibility tree—more specifically this is a disjunctive self-reducibility tree—of a formula, we mean the tree that has the formula as its root, and in which each node corresponding to a formula with some variables unassigned has as its left and right children the same formula but with the lexicographically first unassigned variable set respectively to True and to False.

In the proof of Theorem 1.2, we were greatly helped by the fact that we were dealing with whether tally sets are hard for NP. Tally strings are easily identifiable as such, and that made our pruning scheme straightforward. We now turn to a slightly more difficult case.

Theorem 1.4 *If there is a sparse set that is Σ_1^1 -hard for coNP, then $P = \text{NP}$.*

Corollary 1.5 *If there is a sparse coNP-complete set, then $P = \text{NP}$.*

The proof of Theorem 1.4 goes as follows. As in the proof of Theorem 1.2, we wish to use our hypothesis to construct a polynomial-time algorithm for SAT. Indeed, we wish to do so by expanding and pruning the self-reducibility tree as was done in the proof of Theorem 1.2. The key obstacle is that the pruning procedure from the proof of Theorem 1.2 no longer works, since unlike tally sets, sparse sets are not necessarily "P-capturable" (a set is P-capturable if it is a subset of some sparse P set). In the following proof, we replace the tree-pruning procedure of Theorem 1.2 with a tree-pruning procedure based on the following counting trick. We expand our tree, while pruning only duplicates; we argue that if the tree ever becomes larger than a certain polynomial size, then the very failure of our tree pruning proves that the formula is satisfiable.

Proof of Theorem 1.4 Let S be the (hypothetical) sparse set that is $\leq \frac{n}{m}$ -hard for coNP. For each k , let $p_k(n)$ denote the polynomial $n^k - k$. Let d be such that $(\forall n) \left[|S^{\leq n}| \leq p_d(n) \right]$.² Since $\text{SAT} \in \text{NP}$, it follows that $\overline{\text{SAT}} \in \text{coNP}$. Let g be a deterministic polynomial-time many-one reducing $\overline{\text{SAT}}$ to S . Let k be an integer such that $(\forall x) \left[|g(x)| \leq p_k(n) \right]$; since g is computed by a deterministic polynomial-time Turing machine, such a k indeed exists.

We now give, under the hypothesis of this theorem, a deterministic polynomial-time algorithm for SAT, via a simple tree-pruning procedure. As in the proof of Theorem 1.2, let F be an input formula, and let m be the number of variables in F . Without loss of generality, let $m \geq 1$ and let the variables of F be named v_1, \dots, v_m . Each stage of our construction will pass forward a collection of formulas. View Stage 0 as passing on to the next stage the collection containing just the formula F . We now specify Stage i . Note that Steps 1 and 2 are the same as in the proof of Theorem 1.2, Step 3 is modified, and Step 4 is new.

Stage i , $1 \leq i \leq m$, assuming the collection at the end of Stage $i-1$ is $\{F_1, \dots, F_k\}$.

Step 1 Let C be the collection

$$\{F_1[v_1 = \text{True}], F_2[v_1 = \text{True}], \dots, F_k[v_1 = \text{True}], \\ F_1[v_1 = \text{False}], F_2[v_1 = \text{False}], \dots, F_k[v_1 = \text{False}]\}$$

Step 2 Set C' to be C .

Step 3 For each formula f in C' (in arbitrary order) do:

1. Compute $g(f)$.
2. If for no formula h in C' does $g(f) = g(h)$, then add f to C' .

² The $|S^{\leq n}|$, as opposed to the $|S^n|$ that implicitly appears in the definition of "sparse set" (equation 1.1), is not a typographical error. Both yield valid and equivalent definitions of the class of sparse sets. The $|S^n|$ approach is, as we will see in Chap. 3, a bit more fine-grained. However, the proof of the present theorem works most smoothly with the $|S^{\leq n}|$ definition.

Step 4 If C contains at least $p_0(p_0(|F|)) + 1$ elements, stop and immediately declare that $F \in \text{SAT}$. (The reason for the $p_0(p_0(|F|)) + 1$ figure will be made clear below.)

End Stage i [C' is the collection that gets passed on to Stage $i + 1$].

The action of our algorithm at Stage $m + 1$ is as follows: If some member of the (variable-free) formula collection output by Stage m evaluates to being true we declare $F \in \text{SAT}$, and otherwise we declare $F \notin \text{SAT}$.

Why does this algorithm work? Let α represent $\neg F$. Since $p_0(p_0(n)) + 1$ is a polynomial in the input size, F , it is clear that the above algorithm runs in polynomial time. If the hypothesis of Step 4 is never met, then the algorithm is correct for reasons similar to those showing the correctness of the proof of Theorem 1.2.

If Step 4 is ever invoked, then at the stage at which it is invoked, we have $p_0(p_0(n)) + 1$ distinct strings being mapped to by the non-pruned nodes at the current level of our self-reducibility tree. (Recall that by the self-reducibility tree—more specifically this is a disjunctive self-reducibility tree—of a formula, we mean the tree that has the formula as its root, and in which each node corresponding to a formula with some variables unassigned has as its left and right children the same formula but with the lexicographically first unassigned variable set respectively to True and False.) Note that each of these mapped-to strings is of length at most $p_0(n)$ since that is the longest string that reduction g can output on inputs of size at most n . However, there are only $p_0(p_0(n))$ strings in $S^{S^{\leq p_0(n)}}$. As usual, Σ denotes our alphabet, and as usual we take $\Sigma = \{0, 1\}$. So since the formulas in our collection map to $p_0(p_0(n)) + 1$ distinct strings in $(\Sigma^*)^{S^{\leq p_0(n)}}$, at least one formula in our collection, call it H , maps under the action of g to a string in \bar{S} .³ So $g(H) \notin S$. However, $\overline{\text{SAT}}$ reduces to S via g , so H is satisfiable. Since H was obtained by making substitutions to some variables of F , it follows that F is satisfiable. Thus, if the hypothesis of Step 4 is ever met, it is indeed correct to halt immediately and declare that F is satisfiable. \blacksquare Theorem 1.1

Pause to Ponder 1.6 *In light of the comment in footnote 3, change the proof so that Step 4 does not terminate the algorithm, but rather the algorithm drives forward to explicitly find a satisfying assignment for F . (Hint: The crucial point is to, via pruning, keep the tree from getting too large. The following footnote contains a give-away hint.)*

³ Note that in this case we know that such an H exists, but we have no idea which formula is such an H . See Pause to Ponder 1.6 for how to modify the proof to make it more constructive.

⁴ Change Step 4 so that, as soon as C contains $p_0(p_0(n)) + 1$ formulas, no more elements are added to C at the current level.

1.1.2 The Left Set Technique

1.1.2.1 Sparse Complete Sets for NP. So far we have seen, as the proofs of Theorems 1.2 and 1.4, true priming algorithms that show that “thin” sets cannot be hard for certain complexity classes. Inspired by these two results, Mahaney extended them by proving the following lovely, natural result.

Theorem 1.7 *If NP has sparse complete sets then $P = NP$.*

Pause to Ponder 1.8 *The reader will want to envision him- or herself of the fact that the approach of the proof of Theorem 1.4 utterly fails to establish Theorem 1.7. (See this footnote for why⁵.)*

We will not prove Theorem 1.7 now since we soon prove, as Theorem 1.10, a more general result showcasing the left set technique, and that result will immediately imply Theorem 1.7. Briefly put, the new technique needed to prove Theorems 1.7 and 1.10 is the notion of a “left set.” Very informally, a left set fills in gaps so as to make binary search easier.

Theorem 1.7 establishes that if there is a sparse NP-complete set then $P = NP$. In NP, the existence of sparse NP-hard sets and the existence of sparse NP-complete sets stand or fall together. (One can alternatively conclude this from the fact that Theorem 1.10 establishes its result for NP- \leq_{m}^p -hardness rather than merely for NP- \leq_{m}^p -completeness.)

Theorem 1.9 *NP has sparse \leq_{m}^p -hard sets if and only if NP has sparse \leq_{m}^p -complete sets.*

Proof The “if” direction is immediate. So, we need only prove that if NP has a \leq_{m}^p -hard sparse set then it has a \leq_{m}^p -complete sparse set. Let S be any sparse set that is \leq_{m}^p -hard for NP. Since S is \leq_{m}^p -hard, it holds that $\text{SAT} \leq_{\text{m}}^p S$. Let f be a polynomial-time computable function that many-one reduces SAT to S . Define

$$S' = \{0^k \# y \mid k \geq 0 \wedge (\exists x \in \text{SAT})[k \geq |x| \wedge f(x) = y]\}.$$

The rough intuition here is that S' is almost $f(\text{SAT})$, except to make the proof work it via the 0^k also has a padding part. Note that if $0^k \# z \in S'$ then certainly $z \in S$. S' is clearly in NP, since to test whether $0^k \# z$ is in S' we nondeterministically guess a string x of length at most k and we nondeterministically guess a potential certificate of $x \in \text{SAT}$ (i.e., we guess a complete assignment of the variables of the formula x), and (on each guessed path) we accept if the guessed string/certificate pair is such that $f(x) = z$.

⁵ The analogous proof would merely be able to claim that if the tree were getting “bushy,” there would be at least one *satisfiable* formula among the collection. This says nothing regarding whether some other formula might be satisfiable. Thus, even if the set C' is getting very large, we have no obvious way to prune it.

and the certificate proves that $x \in \text{SAT}$. Given that S is sparse, it is not hard to see that S' is also sparse. Finally, S' is NP-hard because, in light of the fact that $\text{SAT} \leq_m^p S$ via polynomial-time reduction f , it is not hard to see that $\text{SAT} \leq_m^p S'$ via the reduction $f'(x) = 0^{|\alpha|} \# f(x)$. \square

We now turn to presenting the left set technique. We do so via proving that if some sparse set is NP-complete under bounded-truth-table reductions then $P = \text{NP}$.

Theorem 1.10 *If there is a sparse set then $P = \text{NP}$ that is \leq_{btt}^p -hard for NP , then $P = \text{NP}$.*

In the rest of the section, we prove Theorem 1.10.

1.1.2.2 The Left Set and w_{\max} . Let L be an arbitrary member of NP. There exist a polynomial p and a language in $A \in P$ such that, for every $x \in \Sigma^*$,

$$x \in L \iff (\exists w \in \Sigma^{p(|x|)}) [(x, w) \in A].$$

For each $x \in \Sigma^*$ and $w \in \Sigma^*$, call w a *witness* for $x \in L$ with respect to A and p if $|w| = p(|x|)$ and $(x, w) \in A$. Define the *left set* with respect to A and p , denoted by $\text{Left}[A, p]$, to be

$$\{(x, y) \mid x \in \Sigma^* \wedge y \in \Sigma^{p(|x|)} \wedge (\exists w \in \Sigma^{p(|x|)}) [w \geq y \wedge (x, w) \in A]\},$$

i.e., $\text{Left}[A, p]$ is the set of all (x, y) such that y belongs to $\Sigma^{p(|x|)}$ and is “to the left” of some witness for $x \in L$ with respect to A and p . For each $x \in \Sigma^*$, define

$$w_{\max}(x) = \max\{y \in \Sigma^{p(|x|)} \mid (x, y) \in \text{Left}[A, p]\};$$

if $\{y \in \Sigma^{p(|x|)} \mid (x, y) \in A\}$ is empty, then $w_{\max}(x)$ is undefined. In other words, $w_{\max}(x)$ is the lexicographic maximum of the witnesses for $x \in L$ with respect to A and p . Clearly, for every $x \in \Sigma^*$,

$$x \in L \iff w_{\max}(x) \text{ is defined,}$$

and

$$x \in L \iff (\exists y \in \Sigma^{p(|x|)}) [(x, y) \in \text{Left}[A, p]].$$

Furthermore, for every $x \in \Sigma^*$, the set

$$\{y \in \Sigma^{p(|x|)} \mid (x, y) \in \text{Left}[A, p]\}$$

equals $\{y \in \Sigma^{p(|x|)} \mid 0^{p(|x|)} \leq y \leq w_{\max}(x)\}$ if $x \in L$ and equals \emptyset otherwise (see Fig. 1.1). More precisely,

$$(\forall x \in \Sigma^*) (\forall y \in \Sigma^{p(|x|)}) [(x, y) \in \text{Left}[A, p] \iff y \in w_{\max}(x)].$$

Also,

$$\begin{aligned} (\forall x \in \Sigma^*) (\forall y, y' \in \Sigma^{p(|x|)}) [(x, y) \in \text{Left}[A, p] \wedge (y' < y)] \\ \implies (x, y') \in \text{Left}[A, p]. \end{aligned} \quad (1.4)$$




Fig. 1.1 The left set $Left(A, p)$. *Top:* The case when $x \in L$. The arrows above are witnesses for $x \in L$ with respect to A and p , $w_{max}(x)$ is the rightmost arrow, and the shaded area is $\{y \in \Sigma^{p(|x|)} : (x, y) \in Left(A, p)\}$. *Bottom:* The case when $x \notin L$. $w_{max}(x)$ is undefined and $\{y \in \Sigma^{p(|x|)} : (x, y) \in Left(A, p)\} = \emptyset$.

Note that $Left(A, p)$ is in NP via the nondeterministic Turing machine that, on input (x, y) , guesses $w \in \Sigma^{p(|x|)}$, and accepts if

$$(y \in \Sigma^{p(|x|)}) \wedge (y \leq w) \wedge ((x, y) \in A)$$

and rejects otherwise.

Below, we introduce a useful characterization of \leq_{arr}^p reductions. Let $k \geq 1$. A k -truth-table condition is a $(k+1)$ -tuple C such that the first component of C is a boolean function of arity k and each of the other components of C is a member of Σ^* . For a k -truth-table condition $C = (\alpha, v_1, \dots, v_k)$, we call α the truth-table of C , call $\{v_i \mid (\exists! z : 1 \leq z \leq k) (z \neq v_i)\}$ the queries of C , and, for each i , $1 \leq i \leq k$, call v_i the i th query of C . For a language D , we say that the k -truth-table condition $(\alpha, v_1, \dots, v_k)$ is satisfied by D if $\alpha(\chi_D(v_1), \dots, \chi_D(v_k)) = 1$.

Proposition 1.11 Suppose that a language C is \leq_{arr}^p reducible to a language D . Then there exist an integer $k \geq 1$ and a polynomial-time computable function f from Σ^* to the set of all k -truth-table conditions, such that for all $u \in \Sigma^*$,

$$u \in C \iff f(u) \text{ is satisfied by } D.$$

Proof of Proposition 1.11 Suppose that, for some $k > 1$, a language C is \leq_{arr}^p -reducible to a language D via (f_0, B_0) such that $f_0 \in \text{FP}$ and $B_0 \in \text{P}$. For all $u \in \Sigma^*$, there exist some i , $1 \leq i \leq k$, and $v_1, \dots, v_k \in \Sigma^*$, such that

- $f_0(u) = v_1 \# \dots \# v_k \#$, and
- $u \in C \iff v_1 \# \chi_D(v_1) \dots \chi_D(v_k) \in B_0$,

where $\# \notin \Sigma$. Let f_1 be the function from Σ^* to $(\Sigma^* \#)^k$ defined for all $u \in \Sigma^*$ by

$$f_1(u) = v_1 \# \dots \# v_k \#^{k+1-k}.$$

where $f_0(x) = x_1 \# \dots \# x_l \#$. Define B_1 to be the set of all $x \# b_1 \# \dots \# b_k$, $x \in \Sigma^*$, $b_1, \dots, b_k \in \Sigma$, such that

$$x \# b_1 \# \dots \# b_k \in D_C,$$

where l is the integer such that $f_0(x) \in (\Sigma^{\#})^l$. Since $B_1 \in P$ and $f_0 \in FP$, it holds that $D_1 \in P$ and $f_1 \in FP$.

Let $\beta = \chi_D(x)$. Let $u \in \Sigma^*$ and let $f_u(x) = x_1 \# \dots \# x_l \#$. Then

$$u \# \chi_D(x_1) \# \dots \# \chi_D(x_l) \in D_C \iff u \# \chi_D(x_1) \# \dots \# \chi_D(x_l) \beta^{k-l} \in B_1.$$

Since $f_1(x) = x_1 \# \dots \# x_l \# \beta^{k+1-l}$, we have that the pair (f_1, B_1) witnesses that $C \leq_{FP}^1 D$.

Define function f from Σ^* to the set of all k -truth-table conditions as follows: For each $x \in \Sigma^*$,

$$f(x) = (\alpha, v_1, \dots, v_k),$$

where $f_1(x) = x_1 \# \dots \# x_k \#$ and α is the boolean function defined for all $b_1, \dots, b_k \in \{0, 1\}^k$ by

$$\alpha(b_1, \dots, b_k) = \chi_{D_1}(x \# b_1 \# \dots \# b_k).$$

Since $f_1 \in FP$ and k is a constant, $f \in FP$. For every $x \in \Sigma^*$,

$$x \in C \iff x \# \chi_D(x_1) \# \dots \# \chi_D(x_k) \in B$$

and

$$x \# \chi_D(x_1) \# \dots \# \chi_D(x_k) \in B_1 \iff \alpha(\chi_D(x_1), \dots, \chi_D(x_k)) = 1,$$

where $f_1(x) = x_1 \# \dots \# x_k \#$. So, for all $x \in \Sigma^*$,

$$x \in C \iff f(x) \text{ is satisfied by } D$$

Thus, the statement of the proposition holds. \blacksquare Proposition 1.11

Suppose that NP has a sparse \leq_{m}^1 hard set, S . Since S was an undecidable member of NP, it suffices to prove that $L \in P$. Since $\text{Left}[A, p] \in \text{NP}$, $\text{Left}[A, p] \leq_{\text{m}}^1 S$. So, by Proposition 1.11, there exist some $k \geq 1$ and $f \in FP$ such that, for all $x \in \Sigma^*$, $f(x)$ is a k -truth-table condition, and

$$x \in \text{Left}[A, p] \iff f(x) \text{ is satisfied by } S.$$

In preparation for the remainder of the proof, we define some polynomials. Let p_1 be a strictly increasing polynomial such that for all $x \in \Sigma^*$ and $y \in \Sigma^{x+1}$, $|x, y| \leq p_1(|x|)$. Let p_2 be a strictly increasing polynomial such that for all $x \in \Sigma^*$, every query of $f(x)$ has length at most $p_2(|x|)$. Let p_3 be a strictly increasing polynomial such that for all integers $n \geq 0$, $S^{\leq n} \leq p_3(n)$. Define $q(n) = p_3(p_2(p_1(n)))$. Then, for all $x \in \Sigma^*$,

$$||\{w \in S \mid \exists y \in \Sigma^{q(|w|)} \{w \text{ is a query of } f(x, y)\}\}|| \leq q(|w|).$$

Define $r(n) = (k+1)(2q(n)+1)^2$. Also, for each d , $0 \leq d \leq k$, define $v_d(x) = (k-d)2^{d-1}(2q(x)+1)^{k-d}$. Note that $v_0 = r$ and v_k is the constant 1 polynomial.

1.1.2.3 A Polynomial-Time Search Procedure for w_{\max} . To prove that $L \in P$, we will develop a polynomial-time procedure that, on input $x \in \Sigma^*$, generates a list of strings in $\Sigma^{p(|x|)}$ such that, if $w_{\max}(x)$ is defined then the list is guaranteed to contain $w_{\max}(x)$. The existence of such a procedure implies $L \in P$ as follows: Let M be a Turing machine that, on input $x \in \Sigma^*$, runs the enumeration procedure to obtain a list of candidates for $w_{\max}(x)$, and accept if the list contains a string y such that $(x, y) \in A$ and reject otherwise. Since the enumeration procedure runs in polynomial time and $A \in P$, M can be made polynomial-time bounded. Since the output list of the enumeration procedure is guaranteed to contain $w_{\max}(x)$ if it is defined, M accepts if $x \in L$. If $x \notin L$, there is no $y \in \Sigma^{p(|x|)}$ such that $(x, y) \in A$, so M rejects x . Thus, M correctly decides L . Hence, $L \in P$.

In the rest of the proof we will focus on developing such an enumeration procedure. To describe the procedure we need to define some notions.

Let $n \geq 1$ be an integer. Let I be a subset of Σ^n . We say that I is an *interval over Σ^n* if there exist $y, z \in \Sigma^n$ such that

$$y \leq z \text{ and } I = \{v \in \Sigma^n \mid y \leq v \leq z\}.$$

We call y and z respectively the *left end* and the *right end* of I , and write $[y, z]$ to denote I . Let $I = [u, v]$ and $J = [y, z]$ be two intervals over Σ^n . We say that I and J are *disjoint* if they are disjoint as sets, i.e., either $v < y$ or $z < u$. If I and J are disjoint and $v < y$, we say that I is *lexicographically smaller than J* , and write $I < J$.

Let $x \in \Sigma^*$ and let A be a set of pairwise disjoint intervals over Σ^n . We say that A is *nice for x* if

$$x \in L \implies (\exists I \in A)(w_{\max}(x) \in I).$$

Note that for all $x \in \Sigma^*$

- $\{\emptyset, \{0^{p(|x|)}, 1^{p(|x|)}\}\}$ is nice for x regardless of whether $x \in L$, and
- if $x \notin L$, then every set of pairwise disjoint intervals over $\Sigma^{p(|x|)}$ is nice for x .

Let τ be an ordered (possibly empty) list such that, if τ is not empty then each entry of τ is of the form (a, b) for some $a \in \Sigma^n$ and $b \in (0, 1]$. We call such a list a *hypothesis list*. We say that a hypothesis list τ is *correct* if every pair (a, b) in the list satisfies $\chi_S(a) = b$. Let $x \in \Sigma^*$, let A be a set of pairwise disjoint intervals over $\Sigma^{p(|x|)}$, let Γ be a subset of A , and let τ be a hypothesis list. We say that Γ is a *refinement of A for x under τ* if

$$((A \text{ is nice for } x) \wedge (\tau \text{ is correct})) \implies \Gamma \text{ is nice for } x.$$

The following fact states some useful properties of refinements.

- [read Transit pdf, azw \(kindle\), epub, doc, mobi](#)
- [read online American Ghost](#)
- [download Embodied Wisdom: The Collected Papers of Moshe Feldenkrais](#)
- [read **Jasmine and Fire: A Bittersweet Year in Beirut**](#)
- [download Boink: College Sex by the People Having It](#)
- [read online Windows Help & Advice \(November 2015\) pdf, azw \(kindle\), epub](#)

- <http://ramazotti.ru/library/Hatred--The-Psychological-Descent-Into-Violence.pdf>
- <http://pittiger.com/lib/Skinny-Juices--101-Juice-Recipes-for-Detox-and-Weight-Loss.pdf>
- <http://growingsomeroots.com/ebooks/Saturn-and-Melancholy--Studies-in-the-History-of-Natural-Philosophy--Religion--and-Art.pdf>
- <http://academialanguagebar.com/?ebooks/Jasmine-and-Fire--A-Bittersweet-Year-in-Beirut.pdf>
- <http://thermco.pl/library/The-Cultural-Revolution-Cookbook--Simple--Healthy-Recipes-from-China-s-Countryside.pdf>
- <http://fortune-touko.com/library/Stitches--A-Handbook-on-Meaning--Hope-and-Repair.pdf>