

# **Texts in Theoretical Computer Science**

An EATCS Series

Editors: W. Brauer J. Hromkovič G. Rozenberg A. Salomaa

On behalf of the European Association  
for Theoretical Computer Science (EATCS)

---

Advisory Board:

G. Ausiello M. Broy C.S. Calude A. Condon  
D. Harel J. Hartmanis T. Henzinger T. Leighton  
M. Nivat C. Papadimitriou D. Scott

Daniel Kroening · Ofer Strichman

# Decision Procedures

An Algorithmic Point of View

Foreword by Randal E. Bryant

 Springer

Daniel Kroening  
Computing Laboratory  
University of Oxford  
Wolfson Building  
Parks Road  
Oxford, OX1 3QD  
United Kingdom  
daniel.kroening@comlab.ox.ac.uk

Ofer Strichman  
William Davidson Faculty of Industrial  
Engineering and Management  
Technion – Israel Institute of Technology  
Technion City  
Haifa 32000  
Israel  
ofers@ie.technion.ac.il

ISBN 978-3-540-74104-6

e-ISBN 978-3-540-74105-3

Texts in Theoretical Computer Science. An EATCS Series. ISSN 1862-4499

Library of Congress Control Number: 2008924795

ACM Computing Classification (1998): B.5.2, D.2.4, D.2.5, E.1, F.3.1, F.4.1

© 2008 Springer-Verlag Berlin Heidelberg

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 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

The use of general descriptive names, registered 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:* KünkelLopka GmbH, Heidelberg

Printed on acid-free paper

9 8 7 6 5 4 3 2 1

springer.com

---

# Foreword

By **Randal E. Bryant**

Research in decision procedures started several decades ago, but both their practical importance and the underlying technology have progressed rapidly in the last five years. Back in the 1970s, there was a flurry of activity in this area, mostly centered at Stanford and the Stanford Research Institute (SRI), motivated by a desire to apply formal logic to problems in artificial intelligence and software verification. This work laid foundations that are still in use today. Activity dropped off through the 1980s and 90s, accompanied by a general pessimism about automated formal methods. A conventional wisdom arose that computer systems, especially software, were far too complex to reason about formally.

One notable exception to this conventional wisdom was the success of applying Boolean methods to hardware verification, beginning in the early 1990s. Tools such as model checkers demonstrated that useful properties could be proven about industrial scale hardware systems, and that bugs could be detected that had otherwise escaped extensive simulation. These approaches improved on their predecessors by employing more efficient logical reasoning methods, namely ordered binary decision diagrams and Boolean satisfiability solvers. The importance of considering algorithmic efficiency, and even low-level concerns such as cache performance became widely recognized as having a major impact on the size of problems that could be handled.

Representing systems at a detailed Boolean level limited the applicability of early model checkers to control-intensive hardware systems. Trying to model data operations, as well as the data and control structures found in software leads to far too many states, when every bit of a state is viewed as a separate Boolean signal.

One way to raise the level of abstraction for verifying a system is to view data in more abstract terms. Rather than viewing a computer word as a collection of 32 Boolean values, it can be represented as an integer. Rather than viewing a floating point multiplier as a complex collection of Boolean functions, many verification tasks can simply view it as an “uninterpreted

function” computing some repeatable function over its inputs. From this approach came a renewed interest in decision procedures, automating the process of reasoning about different mathematical forms. Some of this work revived methods dating back many years, but alternative approaches also arose that made use of Boolean methods, exploiting the greatly improved performance of Boolean satisfiability (SAT) solvers. Most recently, decision procedures have become quite sophisticated, using the general framework of search-based SAT solvers, integrated with methods for handling the individual mathematical theories.

With the combination of algorithmic improvements and the improved performance of computer systems, modern decision procedures can readily handle problems that far exceed the capacity of their forebearers from the 1970s. This progress has made it possible to apply formal reasoning to both hardware and software in ways that disprove the earlier conventional wisdom. In addition, the many forms of malicious attacks on computer systems have created a program execution environment where seemingly minor bugs can yield serious vulnerabilities, and this has greatly increased the motivation to apply formal methods to software analysis.

Until now, learning the state of the art in decision procedures required assimilating a vast amount of literature, spread across journals and conferences in a variety of different disciplines and over multiple decades. Ideas are scattered throughout these publications, but with no standard terminology or notation. In addition some approaches have been shown to be unsound, and many have proven ineffective. I am therefore pleased that Daniel Kroening and Ofer Strichman have compiled the vast amount of information on decision procedures into a single volume. Enough progress has been made in the field that the results will be of interest to those wishing to apply decision procedures. At the same time, this is a fast moving and active research community, making the work essential reading for the many researchers in the field.


---

## Preface

A *decision procedure* is an algorithm that, given a decision problem, terminates with a correct yes/no answer. In this book, we concentrate on decision procedures for decidable first-order theories that are useful in the context of automated verification and reasoning, theorem proving, compiler optimization, synthesis, and so forth. Since the ability of these techniques to cope with problems arising in industry depends critically on decision procedures, this is a vibrant and prospering research subject for many researchers around the world, both in academia and in industry. Intel and AMD, for example, are developing and using theorem provers and decision procedures as part of their efforts to build circuit verification tools with ever-growing capacity. Microsoft is developing and routinely using decision procedures in several code analysis tools.

Despite the importance of decision procedures, one rarely finds a university course dedicated entirely to this topic; occasionally, it is addressed in courses on algorithms or on logic for computer science. One of the reasons for this situation, we believe, is the lack of a textbook summarizing the main results in the field in an accessible, uniform way. The primary goal of this book is therefore to serve as a textbook for an advanced undergraduate- or graduate-level computer science course. It does not assume specific prior knowledge beyond what is expected from a third-year undergraduate computer science student. The book may also help graduate students entering the field, as currently they are required to gather information from what seems to be an endless list of articles.

The decision procedures that we describe in this book draw from diverse fields such as graph theory, logic, and operations research. These procedures have to be highly efficient, since the problems they solve are inherently hard. They never seem to be efficient enough, however: what we want to be able to prove is always harder than what we *can* prove. Their asymptotic complexity and their performance in practice must always be pushed further. These characteristics are what makes this topic so compelling for research and teaching.



**Fig. 1.** Decision procedures can be rather complex ... those that we consider in this book take formulas of different theories as input, possibly mix them (using the Nelson–Oppen procedure – see Chap. 10), decide their satisfiability (“YES” or “NO”), and, if yes, provide a satisfying assignment

### Which Theories? Which Algorithms?

A first-order theory can be considered “interesting”, at least from a practical perspective, if it fulfills at least these two conditions:

1. The theory is expressive enough to model a real decision problem. Moreover, it is more expressive or more natural for the purpose of expressing some models in comparison with theories that are easier to decide.

2. The theory is either decidable or semidecidable, and more efficiently solvable than theories that are more expressive, at least in practice if not in theory.<sup>1</sup>

All the theories described in this book fulfill these two conditions. Furthermore, they are all used in practice. We illustrate applications of each theory with examples representative of real problems, whether they may be verification of C programs, verification of hardware circuits, or optimizing compilers. Background in any of these problem domains is not assumed, however.

Other than in one chapter, all the theories considered are quantifier-free. The problem of deciding them is NP-complete. In this respect, they can all be seen as “front ends” of any one of them, for example propositional logic. They differ from each other mainly in how naturally they can be used for modeling various decision problems. For example, consider the theory of equality, which we describe in Chaps. 3 and 4: this theory can express any Boolean combination of Boolean variables and expressions of the form  $x_1 = x_2$ , where  $x_1$  and  $x_2$  are variables ranging over, for example, the natural numbers. The problem of satisfying an expression in this theory can be reduced to a satisfiability problem of a propositional logic formula (and vice versa). Hence, there is no difference between propositional logic and the theory of equality in terms of their ability to model decision problems. However, many problems are more naturally modeled with the equality operator and non-Boolean variables.

For each theory that is discussed, there are many alternative decision procedures in the literature. Effort was made to select those procedures that are known to be relatively efficient in practice, and at the same time are based on what we believe to be an interesting idea. In this respect, we cannot claim to have escaped the natural bias that one has towards one’s own line of research.

Every year, new decision procedures and tools are being published, and it is impossible to write a book that reports on this moving target of “the most efficient” decision procedures (the worst-case complexity of most of the competing procedures is the same). Moreover, many of them have never been thoroughly tested against one another. We refer readers who are interested in the latest developments in this field to the SMT-LIB Web page, as well as to the results of the annual tool competition SMT-COMP (see Appendix A). The SMT-COMP competitions are probably the best way to stay up to date as to the relative efficiency of the various procedures and the tools that implement them. One should not forget, however, that it takes much more than a good algorithm to be efficient in practice.

## The Structure and Nature of This Book

The first chapter is dedicated to basic concepts that should be familiar to third- or fourth-year computer science students, such as formal proofs, the

---

<sup>1</sup> Terms such as *expressive* and *decidable* have precise meanings, and are defined in the first chapter.



satisfiability problem, soundness and completeness, and the trade-off between expressiveness and decidability. It also includes the theoretical basis for the rest of the book. From Sect. 1.5 onwards, the chapter is dedicated to more advanced issues that are necessary as a general introduction to the book, and are therefore recommended even for advanced readers. Each of the 10 chapters that follow is mostly self-contained, and generally does not rely on references to other chapters, other than the first introductory chapter. An exception to this rule is Chap. 4, which relies on definitions and explanations given in Chap. 3.

The mathematical symbols and notations are mostly local to each chapter. Each time a new symbol is introduced, it appears in a rounded box in the margin of the page for easy reference. All chapters conclude with problems, varying in level of difficulty, and bibliographic notes and a glossary of symbols.

A draft of this book was used as lecture notes for a combined undergraduate and graduate course on decision procedures at the Technion, Israel, at ETH Zurich, Switzerland, and at Oxford University, UK. The slides that were used in these courses, as well as links to other resources appear on the book's Web page ([www.decision-procedures.org](http://www.decision-procedures.org)). Source code of a C++ library for rapid development of decision procedures can also be downloaded from this page. This library provides the necessary infrastructure for programming many of the algorithms described in this book, as explained in Appendix B. Implementing one of these algorithms was a requirement in the course, and it proved successful. It even led several students to their thesis topic.

## Acknowledgments

Many people read drafts of this manuscript and gave us useful advice. We would like to thank, in alphabetical order, Domagoj Babic, Josh Berdine, Hana Chockler, Leonardo de Moura, Benny Godlin, Alan Hu, Wolfgang Kunz, Shuvendu Lahiri, Albert Oliveras Llunell, Joel Ouaknine, Hendrik Post, Sharon Shoham, Aaron Stump, Cesare Tinelli, Ashish Tiwari, Rachel Tzoref, Helmut Veith, Georg Weissenbacher, and Calogero Zarba. We thank Ilya Yodovsky Jr. for the drawing in Fig. 1.

February 2008

Daniel Kroening  
Oxford University, United Kingdom

Ofer Strichman  
Technion, Haifa, Israel

---

# Contents

<b>1</b>	<b>Introduction and Basic Concepts</b>	1
1.1	Two Approaches to Formal Reasoning	3
1.1.1	Proof by Deduction	3
1.1.2	Proof by Enumeration	4
1.1.3	Deduction and Enumeration	5
1.2	Basic Definitions	5
1.3	Normal Forms and Some of Their Properties	8
1.4	The Theoretical Point of View	14
1.4.1	The Problem We Solve	17
1.4.2	Our Presentation of Theories	17
1.5	Expressiveness vs. Decidability	18
1.6	Boolean Structure in Decision Problems	19
1.7	Problems	21
1.8	Glossary	23
<b>2</b>	<b>Decision Procedures for Propositional Logic</b>	25
2.1	Propositional Logic	25
2.1.1	Motivation	25
2.2	SAT Solvers	27
2.2.1	The Progress of SAT Solving	27
2.2.2	The DPLL Framework	28
2.2.3	BCP and the Implication Graph	30
2.2.4	Conflict Clauses and Resolution	35
2.2.5	Decision Heuristics	39
2.2.6	The Resolution Graph and the Unsatisfiable Core	41
2.2.7	SAT Solvers: Summary	42
2.3	Binary Decision Diagrams	43
2.3.1	From Binary Decision Trees to ROBDDs	43
2.3.2	Building BDDs from Formulas	46
2.4	Problems	50
2.4.1	Warm-up Exercises	50

2.4.2	Modeling . . . . .	50
2.4.3	Complexity . . . . .	51
2.4.4	DPLL SAT Solving . . . . .	52
2.4.5	Related Problems . . . . .	52
2.4.6	Binary Decision Diagrams . . . . .	53
2.5	Bibliographic Notes . . . . .	54
2.6	Glossary . . . . .	57
<b>3</b>	<b>Equality Logic and Uninterpreted Functions . . . . .</b>	<b>59</b>
3.1	Introduction . . . . .	59
3.1.1	Complexity and Expressiveness . . . . .	59
3.1.2	Boolean Variables . . . . .	60
3.1.3	Removing the Constants: A Simplification . . . . .	60
3.2	Uninterpreted Functions . . . . .	60
3.2.1	How Uninterpreted Functions Are Used . . . . .	61
3.2.2	An Example: Proving Equivalence of Programs . . . . .	63
3.3	From Uninterpreted Functions to Equality Logic . . . . .	64
3.3.1	Ackermann's Reduction . . . . .	66
3.3.2	Bryant's Reduction . . . . .	69
3.4	Functional Consistency Is Not Enough . . . . .	72
3.5	Two Examples of the Use of Uninterpreted Functions . . . . .	74
3.5.1	Proving Equivalence of Circuits . . . . .	75
3.5.2	Verifying a Compilation Process with Translation Validation . . . . .	77
3.6	Problems . . . . .	78
3.6.1	Warm-up Exercises . . . . .	78
3.6.2	Problems . . . . .	78
3.7	Glossary . . . . .	79
<b>4</b>	<b>Decision Procedures for Equality Logic and Uninterpreted Functions . . . . .</b>	<b>81</b>
4.1	Congruence Closure . . . . .	81
4.2	Basic Concepts . . . . .	83
4.3	Simplifications of the Formula . . . . .	85
4.4	A Graph-Based Reduction to Propositional Logic . . . . .	88
4.5	Equalities and Small-Domain Instantiations . . . . .	92
4.5.1	Some Simple Bounds . . . . .	93
4.5.2	Graph-Based Domain Allocation . . . . .	94
4.5.3	The Domain Allocation Algorithm . . . . .	96
4.5.4	A Proof of Soundness . . . . .	98
4.5.5	Summary . . . . .	101
4.6	Ackermann's vs. Bryant's Reduction: Where Does It Matter? . . . . .	101
4.7	Problems . . . . .	103
4.7.1	Conjunctions of Equalities and Uninterpreted Functions . . . . .	103
4.7.2	Reductions . . . . .	104

---

4.7.3	Complexity . . . . .	105
4.7.4	Domain Allocation . . . . .	106
4.8	Bibliographic Notes . . . . .	106
4.9	Glossary . . . . .	108
<b>5</b>	<b>Linear Arithmetic</b> . . . . .	<b>111</b>
5.1	Introduction . . . . .	111
5.1.1	Solvers for Linear Arithmetic . . . . .	112
5.2	The Simplex Algorithm . . . . .	113
5.2.1	Decision Problems and Linear Programs . . . . .	113
5.2.2	Basics of the Simplex Algorithm . . . . .	114
5.2.3	Simplex with Upper and Lower Bounds . . . . .	116
5.2.4	Incremental Problems . . . . .	120
5.3	The Branch and Bound Method . . . . .	120
5.3.1	Cutting-Planes . . . . .	122
5.4	Fourier–Motzkin Variable Elimination . . . . .	126
5.4.1	Equality Constraints . . . . .	126
5.4.2	Variable Elimination . . . . .	126
5.4.3	Complexity . . . . .	129
5.5	The Omega Test . . . . .	129
5.5.1	Problem Description . . . . .	129
5.5.2	Equality Constraints . . . . .	130
5.5.3	Inequality Constraints . . . . .	132
5.6	Preprocessing . . . . .	138
5.6.1	Preprocessing of Linear Systems . . . . .	138
5.6.2	Preprocessing of Integer Linear Systems . . . . .	139
5.7	Difference Logic . . . . .	140
5.7.1	Introduction . . . . .	140
5.7.2	A Decision Procedure for Difference Logic . . . . .	142
5.8	Problems . . . . .	142
5.8.1	Warm-up Exercises . . . . .	142
5.8.2	The Simplex Method . . . . .	143
5.8.3	Integer Linear Systems . . . . .	143
5.8.4	Omega Test . . . . .	144
5.8.5	Difference Logic . . . . .	145
5.9	Bibliographic Notes . . . . .	145
5.10	Glossary . . . . .	146
<b>6</b>	<b>Bit Vectors</b> . . . . .	<b>149</b>
6.1	Bit-Vector Arithmetic . . . . .	149
6.1.1	Syntax . . . . .	149
6.1.2	Notation . . . . .	151
6.1.3	Semantics . . . . .	152
6.2	Deciding Bit-Vector Arithmetic with Flattening . . . . .	156
6.2.1	Converting the Skeleton . . . . .	156

6.2.2	Arithmetic Operators	157
6.3	Incremental Bit Flattening	160
6.3.1	Some Operators Are Hard	160
6.3.2	Enforcing Functional Consistency	162
6.4	Using Solvers for Linear Arithmetic	163
6.4.1	Motivation	163
6.4.2	Integer Linear Arithmetic for Bit Vectors	163
6.5	Fixed-Point Arithmetic	165
6.5.1	Semantics	165
6.5.2	Flattening	167
6.6	Problems	167
6.6.1	Semantics	167
6.6.2	Bit-Level Encodings of Bit-Vector Arithmetic	168
6.6.3	Using Solvers for Linear Arithmetic	169
6.7	Bibliographic Notes	169
6.8	Glossary	170
<b>7</b>	<b>Arrays</b>	<b>171</b>
7.1	Introduction	171
7.2	Arrays as Uninterpreted Functions	172
7.3	A Reduction Algorithm for Array Logic	175
7.3.1	Array Properties	175
7.3.2	A Reduction Algorithm	176
7.4	Problems	178
7.5	Bibliographic Notes	178
7.6	Glossary	179
<b>8</b>	<b>Pointer Logic</b>	<b>181</b>
8.1	Introduction	181
8.1.1	Pointers and Their Applications	181
8.1.2	Dynamic Memory Allocation	182
8.1.3	Analysis of Programs with Pointers	184
8.2	A Simple Pointer Logic	185
8.2.1	Syntax	185
8.2.2	Semantics	187
8.2.3	Axiomatization of the Memory Model	188
8.2.4	Adding Structure Types	189
8.3	Modeling Heap-Allocated Data Structures	190
8.3.1	Lists	190
8.3.2	Trees	191
8.4	A Decision Procedure	193
8.4.1	Applying the Semantic Translation	193
8.4.2	Pure Variables	195
8.4.3	Partitioning the Memory	196
8.5	Rule-Based Decision Procedures	197

8.5.1	A Reachability Predicate for Linked Structures . . . . .	198
8.5.2	Deciding Reachability Predicate Formulas . . . . .	199
8.6	Problems . . . . .	202
8.6.1	Pointer Formulas . . . . .	202
8.6.2	Reachability Predicates . . . . .	203
8.7	Bibliographic Notes . . . . .	204
8.8	Glossary . . . . .	206
<b>9</b>	<b>Quantified Formulas</b> . . . . .	<b>207</b>
9.1	Introduction . . . . .	207
9.1.1	Example: Quantified Boolean Formulas . . . . .	209
9.1.2	Example: Quantified Disjunctive Linear Arithmetic . . . . .	211
9.2	Quantifier Elimination . . . . .	211
9.2.1	Prenex Normal Form . . . . .	211
9.2.2	Quantifier Elimination Algorithms . . . . .	213
9.2.3	Quantifier Elimination for Quantified Boolean Formulas . . . . .	214
9.2.4	Quantifier Elimination for Quantified Disjunctive Linear Arithmetic . . . . .	217
9.3	Search-Based Algorithms for QBF . . . . .	218
9.4	Problems . . . . .	220
9.4.1	Warm-up Exercises . . . . .	220
9.4.2	QBF . . . . .	220
9.5	Bibliographic Notes . . . . .	223
9.6	Glossary . . . . .	224
<b>10</b>	<b>Deciding a Combination of Theories</b> . . . . .	<b>225</b>
10.1	Introduction . . . . .	225
10.2	Preliminaries . . . . .	225
10.3	The Nelson–Oppen Combination Procedure . . . . .	227
10.3.1	Combining Convex Theories . . . . .	227
10.3.2	Combining Nonconvex Theories . . . . .	230
10.3.3	Proof of Correctness of the Nelson–Oppen Procedure . . . . .	233
10.4	Problems . . . . .	236
10.5	Bibliographic Notes . . . . .	236
10.6	Glossary . . . . .	239
<b>11</b>	<b>Propositional Encodings</b> . . . . .	<b>241</b>
11.1	Overview . . . . .	241
11.2	Lazy Encodings . . . . .	244
11.2.1	Definitions and Notations . . . . .	244
11.2.2	Building Propositional Encodings . . . . .	245
11.2.3	Integration into DPLL . . . . .	246
11.2.4	Theory Propagation and the DPLL( $T$ ) Framework . . . . .	246
11.2.5	Some Implementation Details of DPLL( $T$ ) . . . . .	250
11.3	Propositional Encodings with Proofs (Advanced) . . . . .	253

11.3.1	Encoding Proofs . . . . .	254
11.3.2	Complete Proofs . . . . .	255
11.3.3	Eager Encodings . . . . .	257
11.3.4	Criteria for Complete Proofs . . . . .	258
11.3.5	Algorithms for Generating Complete Proofs . . . . .	259
11.4	Problems . . . . .	263
11.5	Bibliographic Notes . . . . .	264
11.6	Glossary . . . . .	267
<b>A</b>	<b>The SMT-LIB Initiative . . . . .</b>	<b>269</b>
<b>B</b>	<b>A C++ Library for Developing Decision Procedures . . . . .</b>	<b>271</b>
B.1	Introduction . . . . .	271
B.2	Graphs and Trees . . . . .	272
B.2.1	Adding “Payload” . . . . .	274
B.3	Parsing . . . . .	274
B.3.1	A Grammar for First-Order Logic . . . . .	274
B.3.2	The Problem File Format . . . . .	276
B.3.3	A Class for Storing Identifiers . . . . .	277
B.3.4	The Parse Tree . . . . .	277
B.4	CNF and SAT . . . . .	278
B.4.1	Generating CNF . . . . .	278
B.4.2	Converting the Propositional Skeleton . . . . .	281
B.5	A Template for a Lazy Decision Procedure . . . . .	281
	<b>References . . . . .</b>	<b>285</b>
	<b>Index . . . . .</b>	<b>299</b>

## Introduction and Basic Concepts

While the focus of this book is on algorithms rather than mathematical logic, the two points of view are inevitably mixed: one cannot truly understand why a given algorithm is correct without understanding the logic behind it. This does not mean, however, that logic is a prerequisite, or that without understanding the fundamentals of logic, it is hard to learn and use these algorithms. It is similar, perhaps, to a motorcyclist who has the choice of whether to learn how his or her bike works.

He or she can ride a long way without such knowledge, but at certain points, when things go wrong or if the bike has to be tuned for a particular ride, understanding how and why things work comes in handy. And then again, suppose our motorcyclist does decide to learn mechanics: where should he or she stop? Is the physics of combustion engines important? Is the “why” important at all, or just the “how”? Or an even more fundamental question: should one first learn how to ride a motorcycle and then refer to the basics when necessary, or learn things “bottom-up”, from principles to mechanics – from science to engineering – and then to the rules of driving?

The reality is that different people have different needs, tendencies, and backgrounds, and there is no right way to write a motorcyclist’s manual that fits all. And things can get messier when one is trying to write a book about decision procedures which is targeted, on one hand, at practitioners – programmers who need to know about algorithms that solve their particular problems – and, on the other hand, at students and researchers who need to see how these algorithms can be defined in the theoretical framework that they are accustomed to, namely logic.

This first chapter has been written with both types of reader in mind. It is a combination of a reference for later chapters and a general introduction. Section 1.1 describes the two most common approaches to formal reasoning, namely deduction and enumeration, and demonstrates them with propositional logic. Section 1.2 serves as a reference for basic terminology such as *validity*, *satisfiability*, *soundness* and *completeness*. More basic terminology is described in Sect. 1.3, which is dedicated to *normal forms* and some of their



properties. Up to that point in the chapter, there is no new material. As of Sect. 1.5, the chapter is dedicated to more advanced issues that are necessary as a general introduction to the book. Section 1.4 positions the subject which this book is dedicated to in the theoretical framework in which it is typically discussed in the literature. This is important mainly for the second type of reader: those who are interested in entering this field as researchers, and, more generally, those who are trained to some extent in mathematical logic. This section also includes a description of the types of problem that we are concerned with in this book, and the standard form in which they are presented in the following chapters. Section 1.5 describes the trade-off between expressiveness and decidability. In Sect. 1.6, we conclude the chapter by discussing the need for reasoning about formulas with a Boolean structure.

*What about the rest of the book?* Each chapter is dedicated to a different first-order **theory**. We have not yet explained what a theory is, and specifically what a **first-order theory** is – this is the role of Sect. 1.4 – but some examples are still in order, as some intuition as to what theories are is required before we reach that section in order to understand the direction in which we are proceeding.

Informally, one may think of a theory as a finite or an infinite set of formulas, which are characterized by common grammatical rules, allowed functions and predicates, and a domain of values. The fact that they are called “first-order” means only that there is a restriction on the quantifiers (only variables, rather than sets of variables, can be quantified), but this is mostly irrelevant to us, because, in all chapters but one, we restrict the discussion to quantifier-free formulas. The table below lists some of the first-order theories that are covered in this book.<sup>1</sup>

Theory name	Example formula	Chapter
Propositional logic	$x_1 \wedge (x_2 \vee \neg x_3)$	2
Equality	$y_1 = y_2 \wedge \neg(y_1 = y_3) \implies \neg(y_1 = y_3)$	3, 4
Linear arithmetic	$(2z_1 + 3z_2 \leq 5) \vee (z_2 + 5z_2 - 10z_3 \geq 6)$	5
Bit vectors	$((a \gg b) \& c) < c$	6
Arrays	$(i = j \wedge a[j] = 1) \implies a[i] = 1$	7
Pointer logic	$p = q \wedge *p = 5 \implies *q = 5$	8
Combined theories	$(i \leq j \wedge a[j] = 1) \implies a[i] < 2$	10

In the next few sections, we use propositional logic, which we assume the reader is familiar with, in order to demonstrate various concepts that apply equally to other first-order theories.

<sup>1</sup> Here we consider propositional logic as a first-order theory, which is technically correct, although not common.

## 1.1 Two Approaches to Formal Reasoning

The primary problem that we are concerned with is that of the validity (or satisfiability) of a given formula. Two fundamental strategies for solving this problem are the following:

- The **model-theoretic** approach is to enumerate possible solutions from a finite number of candidates.
- The **proof-theoretic** approach is to use a deductive mechanism of reasoning, based on **axioms** and **inference rules**, which together are called an **inference system**.

These two directions – enumeration and deduction – are apparent as early as the first lessons on propositional logic. We dedicate this section to demonstrating them.

Consider the following three contradicting claims:

1. If  $x$  is a prime number greater than 2, then  $x$  is odd.
2. It is not the case that  $x$  is not a prime number greater than 2.
3.  $x$  is not odd.

Denote the statement “ $x$  is a prime number greater than 2” by  $A$  and the statement “ $x$  is odd” by  $B$ . These claims translate into the following propositional formulas:

$$\begin{aligned} A &\implies B . \\ \neg\neg A &. \\ \neg B &. \end{aligned} \tag{1.1}$$

We would now like to prove that this set of formulas is indeed inconsistent.

### 1.1.1 Proof by Deduction

The first approach is to derive conclusions by using an inference system. Inference rules relate **antecedents** to their **consequents**. For example, the following are two inference rules, called modus ponens (M.P.) and CONTRADICTION:

$$\frac{\varphi_1 \implies \varphi_2 \quad \varphi_1}{\varphi_2} \text{ (M.P.)} , \tag{1.2}$$

$$\frac{\varphi \quad \neg\varphi}{\text{FALSE}} \text{ (CONTRADICTION)} . \tag{1.3}$$

The rule M.P. can be read as follows: from  $\varphi_1 \implies \varphi_2$  and  $\varphi_1$  being TRUE, deduce that  $\varphi_2$  is TRUE. The formula  $\varphi_2$  is the consequent of the rule M.P. Axioms are inference rules without antecedents:

$$\frac{}{\neg\neg\varphi \iff \varphi} \text{ (DOUBLE-NEGATION-AX)} . \tag{1.4}$$

(Axioms are typically written without the separating line above them.) We can also write a similar inference rule:

$$\frac{\neg\neg\varphi}{\varphi} \text{ (DOUBLE-NEGATION) .} \quad (1.5)$$

(DOUBLE-NEGATION-AX and DOUBLE-NEGATION are not the same, because the latter is not symmetric.) Many times, however, axioms and inference rules are interchangeable, so there is not always a sharp distinction between them.

The inference rules and axioms above are expressed with the help of arbitrary formula symbols (such as  $\varphi_1$  and  $\varphi_2$  in (1.2)). In order to use them for proving a particular theorem, they need to be *instantiated*, which means that these arbitrary symbols are replaced with specific variables and formulas that are relevant to the theorem that we wish to prove. For example, the inference rules (1.2), (1.3), and (1.5) can be instantiated such that FALSE, i.e., a contradiction, can be derived from the set of formulas in (1.1):

$$\begin{array}{ll} (1) & A \implies B \text{ (premise)} \\ (2) & \neg\neg A \quad \text{(premise)} \\ (3) & A \quad \text{(2; DOUBLE-NEGATION)} \\ (4) & \neg B \quad \text{(premise)} \\ (5) & B \quad \text{(1, 3; M.P.)} \\ (6) & \text{FALSE} \quad \text{(4, 5; CONTRADICTION) .} \end{array} \quad (1.6)$$

Here, in step (3),  $\varphi$  in the rule DOUBLE-NEGATION is instantiated with  $A$ . The antecedent  $\varphi_1$  in the rule M.P. is instantiated with  $A$ , and  $\varphi_2$  is instantiated with  $B$ .

More complicated theorems may require more complicated inference systems. This raises the question of whether everything that can be proven with a given inference system is indeed valid (in this case the system is called **sound**), and whether there exists a proof of validity using the inference system for every valid formula (in this case it is called **complete**). These questions are fundamental for every deduction system; we delay further discussion of this subject and a more precise definition of these terms to Sect. 1.2.

While deductive methods are very general, they are not always the most convenient or the most efficient way to know whether a given formula is valid.

### 1.1.2 Proof by Enumeration

The second approach is relevant if the problem of checking whether a formula is satisfiable can be reduced to a problem of *searching* for a satisfying assignment within a finite set of options. This is the case, for example, if the variables range over a finite domain,<sup>2</sup> such as in propositional logic. In the case of propositional logic, enumerating solutions can be done using **truth tables**, as demonstrated by the following example:

<sup>2</sup> A finite domain is a sufficient but not a necessary condition. In many cases, even if the domain is infinite, it is possible to find a bound such that if there exists a satisfying assignment, then there exists one within this bound. Theories that have this property are said to have the **small-model property**.

$A$	$B$	$(A \implies B)$	$(A \implies B) \wedge A$	$(A \implies B) \wedge A \wedge \neg B$
1	1	1	1	0
1	0	0	0	0
0	1	1	0	0
0	0	1	0	0

The rightmost column, which represents the formula in our example (see (1.1)), is not satisfied by any one of the four possible assignments, as expected.

### 1.1.3 Deduction and Enumeration

The two basic approaches demonstrated above, deduction and enumeration, go a long way, and in fact are major subjects in the study of logic. In practice, many decision procedures are not based on explicit use of either enumeration or deduction. Yet, typically their actions can be understood as performing one or the other (or both) implicitly, which is particularly helpful when arguing for their correctness.

## 1.2 Basic Definitions

We begin with several basic definitions that are used throughout the book. Some of the definitions that follow do not fully coincide with those that are common in the study of mathematical logic. The reason for these gaps is that we focus on quantifier-free formulas, which enables us to simplify various definitions. We discuss these issues further in Sect. 1.4.

**Definition 1.1 (assignment).** *Given a formula  $\varphi$ , an assignment of  $\varphi$  from a domain  $D$  is a function mapping  $\varphi$ 's variables to elements of  $D$ . An assignment to  $\varphi$  is full if all of  $\varphi$ 's variables are assigned, and partial otherwise.*

In this definition, we assume that there is a single domain for all variables. The definition can be trivially extended to the case in which different variables have different domains.

**Definition 1.2 (satisfiability, validity and contradiction).** *A formula is satisfiable if there exists an assignment of its variables under which the formula evaluates to TRUE. A formula is a contradiction if it is not satisfiable. A formula is valid (also called a tautology) if it evaluates to TRUE under all assignments.*

What does it mean that a formula “evaluates to TRUE” under an assignment? To evaluate a formula, one needs a definition of the **semantics** of the various functions and predicates in the formula. In propositional logic, for example, the semantics of the propositional connectives is given by truth tables, as presented above. Indeed, given an assignment of all variables in a propositional

formula, a truth table can be used for checking whether it satisfies a given formula, or, in other words, whether the given formula evaluates to TRUE under this assignment.

It is not hard to see that a formula  $\varphi$  is valid if and only if  $\neg\varphi$  is a contradiction. Although somewhat trivial, this is a very useful observation, because it means that we can check whether a formula is valid by checking instead whether its negation is a contradiction, i.e., not satisfiable.

**Example 1.3.** The propositional formula

$$A \wedge B \tag{1.7}$$

is satisfiable because there exists an assignment, namely  $\{A \mapsto \text{TRUE}, B \mapsto \text{TRUE}\}$ , which makes the formula evaluate to TRUE. The formula

$$(A \implies B) \wedge A \wedge \neg B \tag{1.8}$$

is a contradiction, as we saw earlier: no assignment satisfies it. On the other hand, the negation of this formula, i.e.,

$$\neg((A \implies B) \wedge A \wedge \neg B), \tag{1.9}$$

is valid: every assignment satisfies it. ▀

Given a formula  $\varphi$  and an assignment  $\alpha$  of its variables, we write  $\alpha \models \varphi$  to denote that  $\alpha$  satisfies  $\varphi$ . If a formula  $\varphi$  is valid (and hence, all assignments satisfy it), we write  $\models \varphi$ .<sup>3</sup>

**Definition 1.4 (the decision problem for formulas).** *The decision problem for a given formula  $\varphi$  is to determine whether  $\varphi$  is valid.*

Given a theory  $T$ , we are interested in a procedure<sup>4</sup> that terminates with a correct answer to the decision problem, for every formula of the theory  $T$ .<sup>5</sup>

This can be formalized with a generalization of the notions of “soundness” and “completeness” that we saw earlier in the context of inference systems. These terms can be defined for the more general case of procedures as follows:

<sup>3</sup> Recall that the discussion here refers to propositional logic. In the more general case, we are not talking about assignments, rather about structures that may or may not satisfy a formula. In that case, the notation  $\models \varphi$  means that all structures satisfy  $\varphi$ . These terms are explained later in Sect. 1.4.

<sup>4</sup> We follow the convention by which a **procedure** does not necessarily terminate, whereas an **algorithm** terminates. This may cause confusion, because a “decision procedure” is by definition terminating, and thus should actually be called a “decision algorithm”. This confusion is rooted in the literature, and we follow it here.

<sup>5</sup> Every theory is defined over a set of symbols (e.g., linear arithmetic is defined over symbols such as “+” and “≥”). By saying “every formula of the theory” we mean every formula that is restricted to the symbols of the theory. This will be explained in more detail in Sect. 1.4.

**Definition 1.5 (soundness of a procedure).** *A procedure for the decision problem is sound if when it returns “Valid”, the input formula is valid.*

**Definition 1.6 (completeness of a procedure).** *A procedure for the decision problem is complete if*

- *it always terminates, and*
- *it returns “Valid” when the input formula is valid.*

**Definition 1.7 (decision procedure).** *A procedure is called a decision procedure for  $T$  if it is sound and complete with respect to every formula of  $T$ .*

**Definition 1.8 (decidability of a theory).** *A theory is decidable if and only if there is a decision procedure for it.*

Given these definitions, we are able to classify procedures according to whether they are sound and complete or only sound. It is rarely the case that unsound procedures are of interest. Ideally, we would always like to have a decision procedure, as defined above. However, sometimes either this is not possible (if the problem is undecidable) or the problem is easier to solve with an incomplete procedure. Some incomplete procedures are categorized as such because they do not *always* terminate (or they terminate with a “don’t know” answer). However, in many practical cases, they do terminate. Thus, completeness can also be thought of as a quantitative property rather than a binary one.

All the theories that we consider in this book are decidable. Once a theory is decidable, the next question is how difficult it is to decide it. A common measure is that of the worst-case or average-case complexity, parameterized by certain characteristics of the input formula, for example its size. One should distinguish between the complexity of a problem and the complexity of an algorithm. For example, most of the decision problems that we consider in this book are in the same complexity class, namely they are NP-complete, but we present different algorithms with different worst-case complexities to solve them. Moreover, since the worst-case complexities of alternative algorithms are frequently the same, we take a pragmatic point of view: is a given decision procedure faster than its alternatives *on a significant set of real benchmark formulas?*

Comparing decision procedures with the same worst-case complexity is problematic: it is rare that one procedure dominates another. The common practice is to consider a decision procedure relevant if it is able to perform faster than others on some significant subset of public benchmarks, or on some well-defined subclass of problems. When there is no way to predict the relative performance of procedures without actually running them, they can be run in parallel, with a “first-to-end kills all others” policy. This is a common practice in industry.

### 1.3 Normal Forms and Some of Their Properties

The term **normal form**, in the context of formulas, is commonly used to indicate that a formula has certain syntactic properties. In this chapter, we introduce normal forms that refer to the Boolean structure of the formula. It is common to begin the process of deciding whether a given formula is satisfiable by transforming it to some normal form that the decision procedure is designed to work with. In order to argue that the overall procedure is correct, we need to show that the transformation preserves satisfiability. The relevant term for describing this relation is the following.

**Definition 1.9 (equisatisfiability).** *Two formulas are equisatisfiable if they are both satisfiable or they are both unsatisfiable.*

The basic blocks of a first-order formula are its predicates, also called the **atoms** of the formula. For example, Boolean variables are the atoms of propositional logic, whereas equalities of the form  $x_i = x_j$  are the atoms of the theory of equality that is studied in Chap. 4.

**Definition 1.10 (negation normal form (NNF)).** *A formula is in negation normal form (NNF) if negation is allowed only over atoms, and  $\wedge, \vee, \neg$  are the only allowed Boolean connectives.*

For example,  $\neg(x_1 \vee x_2)$  is *not* an NNF formula, because the negation is applied to a subformula which is not an atom.

Every quantifier-free formula with a Boolean structure can be transformed in linear time to NNF, by rewriting  $\implies$ ,

$$(a \implies b) \equiv (\neg a \vee b), \quad (1.10)$$

and applying repeatedly what are known as **De Morgan's rules**,

$$\begin{aligned} \neg(a \vee b) &\equiv (\neg a \wedge \neg b), \\ \neg(a \wedge b) &\equiv (\neg a \vee \neg b). \end{aligned} \quad (1.11)$$

In the case of the formula above, this results in  $\neg x_1 \wedge \neg x_2$ .

**Definition 1.11 (literal).** *A literal is either an atom or its negation. We say that a literal is negative if it is a negated atom, and positive otherwise.*

For example, in the propositional-logic formula

$$(a \vee \neg b) \wedge \neg c, \quad (1.12)$$

the set of literals is  $\{a, \neg b, \neg c\}$ , where the last two are negative. In the theory of equality, where the atoms are equality predicates, a set of literals can be  $\{x_1 = x_2, \neg(x_1 = x_3), \neg(x_2 = x_1)\}$ .

Literals are syntactic objects. The set of literals of a given formula changes if we transform it by applying De Morgan's rules. Formula (1.12), for example, can be written as  $\neg(\neg a \wedge b) \wedge \neg c$ , which changes its set of literals.

**Definition 1.12 (state of a literal under an assignment).** A positive literal is satisfied if its atom is assigned TRUE. Similarly, a negative literal is satisfied if its atom is assigned FALSE.

**Definition 1.13 (pure literal).** A literal is called pure in a formula  $\varphi$ , if all occurrences of its variable have the same sign.

In many cases, it is necessary to refer to the set of a formula's literals as if this formula were in NNF. In such cases, either it is assumed that the input formula is in NNF (or transformed to NNF as a first step), or the set of literals in this form is computed indirectly. This can be done by simply counting the number of negations that nest each atom instance: it is negative if and only if this number is odd.

For example,  $\neg x_1$  is a literal in the NNF of


$$\varphi := \neg(\neg x_1 \implies x_2), \quad (1.13)$$

because there is an occurrence of  $x_1$  in  $\varphi$  that is nested in three negations (the fact that  $x_1$  is on the left-hand side of an implication is counted as a negation). It is common in this case to say that the **polarity** (also called the **phase**) of this occurrence is negative.

**Theorem 1.14 (monotonicity of NNF).** Let  $\varphi$  be a formula in NNF and let  $\alpha$  be an assignment of its variables. Let the positive set of  $\alpha$  with respect to  $\varphi$ , denoted  $\text{pos}(\alpha, \varphi)$ , be the literals that are satisfied by  $\alpha$ . For every assignment  $\alpha'$  to  $\varphi$ 's variables such that  $\text{pos}(\alpha, \varphi) \subseteq \text{pos}(\alpha', \varphi)$ ,  $\alpha \models \varphi \implies \alpha' \models \varphi$ .

pos

Figure 1.1 illustrates this theorem: increasing the set of literals satisfied by an assignment maintains satisfiability. It does *not* maintain unsatisfiability, however: it can turn an unsatisfying assignment into a satisfying one.



**Fig. 1.1.** Illustration of Theorem 1.14. The ellipses correspond to the sets of literals satisfied by  $\alpha$  and  $\alpha'$ , respectively

The proof of this theorem is left as an exercise (Problem 1.3).

**Example 1.15.** Let

$$\varphi := (\neg x \wedge y) \vee z \quad (1.14)$$

be an NNF formula. Consider the following assignments and their corresponding positive sets with respect to  $\varphi$ :



$$\begin{aligned} \alpha &:= \{x \mapsto 0, y \mapsto 1, z \mapsto 0\} & \text{pos}(\alpha, \varphi) &:= \{\neg x, y\} \\ \alpha' &:= \{x \mapsto 0, y \mapsto 1, z \mapsto 1\} & \text{pos}(\alpha', \varphi) &:= \{\neg x, y, z\}. \end{aligned} \quad (1.15)$$

By Theorem 1.14, since  $\alpha \models \varphi$  and  $\text{pos}(\alpha, \varphi) \subseteq \text{pos}(\alpha', \varphi)$ , then  $\alpha' \models \varphi$ . Indeed,  $\alpha' \models \varphi$ .  $\blacksquare$

We now describe two very useful restrictions of NNF: disjunctive normal form (DNF) and conjunctive normal form (CNF).

**Definition 1.16 (disjunctive normal form (DNF)).** *A formula is in disjunctive normal form if it is a disjunction of conjunctions of literals, i.e., a formula of the form*

$$\bigvee_i \left( \bigwedge_j l_{ij} \right), \quad (1.16)$$

where  $l_{ij}$  is the  $j$ -th literal in the  $i$ -th **term** (a term is a conjunction of literals).

**Example 1.17.** In propositional logic,  $l$  is a Boolean literal, i.e., a Boolean variable or its negation. Thus the following formula over Boolean variables  $a$ ,  $b$ ,  $c$ , and  $d$  is in DNF:

$$\begin{aligned} (a \wedge c \wedge \neg b) &\vee \\ (\neg a \wedge d) &\vee \\ (b \wedge \neg c \wedge \neg d) &\vee \\ &\vdots \end{aligned} \quad (1.17)$$

In the theory of equality, the atoms are equality predicates. Thus, the following formula is in DNF:

$$\begin{aligned} ((x_1 = x_2) \wedge \neg(x_2 = x_3) \wedge \neg(x_3 = x_1)) &\vee \\ (\neg(x_1 = x_4) \wedge (x_4 = x_2)) &\vee \\ ((x_2 = x_3) \wedge \neg(x_3 = x_4) \wedge \neg(x_4 = x_1)) &\vee \\ &\vdots \end{aligned} \quad (1.18)$$

$\blacksquare$

Every formula with a Boolean structure can be transformed into DNF, while potentially increasing the size of the formula exponentially. The following example demonstrates this exponential ratio.

**Example 1.18.** The following formula is of length linear in  $n$ :

$$(x_1 \vee x_2) \wedge \cdots \wedge (x_{2n-1} \vee x_{2n}). \quad (1.19)$$

The length of the equivalent DNF, however, is exponential in  $n$ , since every new binary clause (a disjunction of two literals) doubles the number of terms in the equivalent DNF, resulting, overall, in  $2^n$  terms:

$$\begin{aligned}
& (x_1 \wedge x_3 \wedge \cdots \wedge x_{2n-3} \wedge x_{2n-1}) \vee \\
& (x_1 \wedge x_3 \wedge \cdots \wedge x_{2n-3} \wedge x_{2n}) \vee \\
& (x_1 \wedge x_3 \wedge \cdots \wedge x_{2n-2} \wedge x_{2n}) \vee \\
& \vdots
\end{aligned} \tag{1.20}$$

▀

Although transforming a formula to DNF can be too costly in terms of computation time, it is a very natural way to decide formulas with an arbitrary Boolean structure.

Suppose we are given a disjunctive linear arithmetic formula, that is, a Boolean structure in which the atoms are linear inequalities over the reals. We know how to decide whether a *conjunction* of such literals is satisfiable: there is a known method called simplex that can give us this answer. In order to use the simplex method to solve the more general case in which there are also disjunctions in the formula, we can perform **syntactic case-splitting**. This means that the formula is transformed into DNF, and then each term is solved separately. Each such term contains a conjunction of literals, a form which we know how to solve. The overall formula is satisfiable, of course, if any one of the terms is satisfiable. **Semantic case-splitting**, on the other hand, refers to techniques that split the search space, in the case where the variables are finite (“first the case in which  $x = 0$ , then the case in which  $x = 1 \dots$ ”).

The term **case-splitting** (without being prefixed with “syntactic”) usually refers in the literature to either syntactic case-splitting or a “smart” implementation thereof. Indeed, many of the cases that are generated in syntactic case-splitting are redundant, i.e., they share a common subset of conjuncts that contradict each other. Efficient decision procedures should somehow avoid replicating the process of deducing this inconsistency, or, in other words, they should be able to **learn**, as demonstrated in the following example.

**Example 1.19.** Consider the following formula:

$$\varphi := (a = 1 \vee a = 2) \wedge a \geq 3 \wedge (b \geq 4 \vee b \leq 0). \tag{1.21}$$

The DNF of  $\varphi$  consists of four terms:

$$\begin{aligned}
& (a = 1 \wedge a \geq 3 \wedge b \geq 4) \vee \\
& (a = 2 \wedge a \geq 3 \wedge b \geq 4) \vee \\
& (a = 1 \wedge a \geq 3 \wedge b \leq 0) \vee \\
& (a = 2 \wedge a \geq 3 \wedge b \leq 0).
\end{aligned} \tag{1.22}$$

These four cases can each be discharged separately, by using a decision procedure for linear arithmetic (Chap. 5). However, observe that the first and the third case share the two conjuncts  $a = 1$  and  $a \geq 3$ , which already makes the case unsatisfiable. Similarly, the second and the fourth case share the conjuncts  $a = 2$  and  $a \geq 3$ . Thus, with the right learning mechanism, two of the

four calls to the decision procedure can be avoided. This is still case-splitting, but more efficient than a plain transformation to DNF.  $\blacksquare$

The problem of reasoning about formulas with a general Boolean structure is a common thread throughout this book.

**Definition 1.20 (conjunctive normal form (CNF)).** *A formula is in conjunctive normal form if it is a conjunction of disjunctions of literals, i.e., it has the form*

$$\bigwedge_i \left( \bigvee_j l_{ij} \right), \quad (1.23)$$

where  $l_{ij}$  is the  $j$ -th literal in the  $i$ -th **clause** (a clause is a disjunction of literals).

Every formula with a Boolean structure can be transformed into an equivalent CNF formula, while potentially increasing the size of the formula exponentially. Yet, any propositional formula can also be transformed into an equisatisfiable CNF formula with only a *linear* increase in the size of the formula. The price to be paid is  $n$  new Boolean variables, where  $n$  is the number of **logical gates** in the formula. This transformation is done via **Tseitin's encoding** [195].

Tseitin suggested that one new variable should be added for every *logical gate* in the original formula, and several clauses to constrain the value of this variable to be equal to the gate it represents, in terms of the inputs to this gate. The original formula is satisfiable if and only if the conjunction of these clauses together with the new variable associated with the topmost operator is satisfiable. This is best illustrated with an example.

**Example 1.21.** Given a propositional formula


$$x_1 \implies (x_2 \wedge x_3), \quad (1.24)$$

with Tseitin's encoding we assign a new variable to each subexpression, or, in other words, to each logical gate, for example AND ( $\wedge$ ), OR ( $\vee$ ), and NOT ( $\neg$ ). For this example, let us assign the variable  $a_2$  to the AND gate (corresponding to the subexpression  $x_2 \wedge x_3$ ) and  $a_1$  to the IMPLICATION gate (corresponding to  $x_1 \implies a_2$ ), which is also the topmost operator of this formula. Figure 1.2 illustrates the **derivation tree** of our formula, together with these auxiliary variables in square brackets.

We need to satisfy  $a_1$ , together with two equivalences,

$$\begin{aligned} a_1 &\iff (x_1 \implies a_2), \\ a_2 &\iff (x_2 \wedge x_3). \end{aligned} \quad (1.25)$$

The first equivalence can be rewritten in CNF as



**Fig. 1.2.** Tseitin's encoding. Assigning an auxiliary variable to each logical gate (shown here in square brackets) enables us to translate each propositional formula to CNF, while increasing the size of the formula only linearly

$$\begin{aligned}
 (a_1 \vee x_1) & \quad \wedge \\
 (a_1 \vee \neg a_2) & \quad \wedge \\
 (\neg a_1 \vee \neg x_1 \vee a_2), &
 \end{aligned} \tag{1.26}$$

and the second equivalence can be rewritten in CNF as

$$\begin{aligned}
 (\neg a_2 \vee x_2) & \quad \wedge \\
 (\neg a_2 \vee x_3) & \quad \wedge \\
 (a_2 \vee \neg x_2 \vee \neg x_3). &
 \end{aligned} \tag{1.27}$$

Thus, the overall CNF formula is the conjunction of (1.26), (1.27), and the unit clause

$$(a_1), \tag{1.28}$$

which represents the topmost operator. ▀

There are various optimizations that can be performed in order to reduce the size of the resulting formula and the number of additional variables. For example, consider the following formula:

$$x_1 \vee (x_2 \wedge x_3 \wedge x_4 \wedge x_5). \tag{1.29}$$

With Tseitin's encoding, we need to introduce four auxiliary variables. The encoding of the clause on the right-hand side, however, can be optimized to use just a single variable, say  $a_2$ :

$$a_2 \iff (x_2 \wedge x_3 \wedge x_4 \wedge x_5). \tag{1.30}$$

In CNF,

$$\begin{aligned}
 (\neg a_2 \vee x_2) & \quad \wedge \\
 (\neg a_2 \vee x_3) & \quad \wedge \\
 (\neg a_2 \vee x_4) & \quad \wedge \\
 (\neg a_2 \vee x_5) & \quad \wedge \\
 (a_2 \vee \neg x_2 \vee \neg x_3 \vee \neg x_4 \vee \neg x_5). &
 \end{aligned} \tag{1.31}$$

In general, we can encode a conjunction of  $n$  literals with a single variable and  $n + 1$  clauses, which is an improvement over the original encoding, requiring  $n - 1$  auxiliary variables and  $3(n - 1)$  clauses.

Such savings are also possible for a series of disjunctions (see Problem 1.1). Another popular optimization is that of **subsumption**: given two clauses such that the set of literals in one of the clauses subsumes the set of literals in the other clause, the longer clause can be discarded without affecting the satisfiability of the formula.

Finally, if the original formula is in NNF, the number of clauses can be reduced substantially, as was shown by Plaisted and Greenbaum in [152]. Tseitin’s encoding is based on constraints of the form

$$\text{auxiliary variable} \iff \text{formula}, \quad (1.32)$$

but only the left-to-right implication is necessary. The proof that this improvement is correct is left as an exercise (Problem 1.4). In practice, experiments show that owing to the requirement of transforming the formula to NNF first, this reduction has a relatively small (positive) effect on the run time of modern SAT solvers compared with Tseitin’s encoding.

**Example 1.22.** Consider a gate  $x_1 \wedge x_2$ , which we encode with a new auxiliary variable  $a$ . Three clauses are necessary to encode the constraint  $a \iff (x_1 \wedge x_2)$ , as was demonstrated in (1.27). The constraint  $a \iff (x_1 \wedge x_2)$  (equivalently,  $(a \vee \neg x_1 \vee \neg x_2)$ ) is redundant, however, which means that only two out of the three constraints are necessary.  $\blacksquare$

A conversion algorithm with similar results to [152], in which the elimination of the negations is built in (rather than the formula being converted to NNF a priori), has been given by Wilson [201].

## 1.4 The Theoretical Point of View

While we take the algorithmic point of view in this book, it is important to understand also the theoretical context, especially for readers who are also interested in following the literature in this field or are more used to the terminology of formal logic. It is also necessary for understanding Chaps. 10 and 11. We must assume in this subsection that the reader is familiar to some extent with first-order logic – a reasonable exposition of this subject is beyond the scope of this book. See [30, 91] for a more organized study of these matters. Let us recall some of the terms that are directly relevant to our topic.

**First-order logic** (also called **predicate logic**) is based on the following elements:

1. *Variables*: a set of *variables*.
2. *Logical symbols*: the standard Boolean connectives (e.g., “ $\wedge$ ”, “ $\neg$ ”, and “ $\vee$ ”), quantifiers (“ $\exists$ ” and “ $\forall$ ”) and parentheses.
3. *Nonlogical symbols*: function, predicate, and constant symbols.
4. *Syntax*: rules for constructing formulas. Formulas adhering to these rules are said to be **well-formed**.

Essentially, first-order logic extends propositional logic with quantifiers and the nonlogical symbols. The syntax of first-order logic extends the syntax of propositional logic naturally. Two examples of such formulas are

- $\exists y \in \mathbb{Z}. \forall x \in \mathbb{Z}. x > y$ ,
- $\forall n \in \mathbb{N}. \exists p \in \mathbb{N}. n > 1 \implies (isprime(p) \wedge n < p < 2n)$ ,

where “ $>$ ”, “ $<$ ”, and “*isprime*” are nonlogical binary predicate symbols.

The elements listed above only refer to symbols and syntax – they still do not tell us how to evaluate whether a given formula is true or false. This separation between symbols and their interpretation – between syntax and semantics – is an important principle in the study of logic. We shall explain this separation with an example. Let  $\Sigma$  denote the set of symbols  $\{0, 1, +, =\}$ , where “0” and “1” are constant symbols, “+” is a binary function symbol, and “=” is a binary predicate symbol. Consider the following formula over  $\Sigma$ :

$$\varphi := \exists x. x + 0 = 1. \quad (1.33)$$

Now, is  $\varphi$  true in  $\mathbb{N}_0$ ? ( $\mathbb{N}_0$  denotes the naturals, including 0.)

What seems like a trivial question is not so simple in the world of formal logic. A logician would say that the answer depends, among other things, on the **interpretation** of the symbols in  $\Sigma$ . What does the “+” symbol mean? Which elements in the domain do “0” and “1” refer to? From a formal perspective, whether  $\varphi$  is true can only be answered with respect to a given **structure**. A structure is a tuple consisting of

- a domain;
- an interpretation of the nonlogical symbols, in the form of a mapping between each function and predicate symbol to a function and a predicate, respectively, and an assignment of a domain element to each of the constant symbols;
- an assignment of a domain element to each of the free (unquantified) variables.

For example, if we choose to interpret the “+” symbol as the *multiplication* function, the answer is that  $\varphi$  in (1.33) is false.

The formula  $\varphi$  is **satisfiable** if and only if there *exists* a structure under which the formula is true. Indeed, in this case there exists such a domain and interpretation – namely,  $\mathbb{N}_0$  and the common interpretation of “+”, “=”, “0” and “1” – and, hence, the formula is satisfiable.

First-order logic can be thought of as a framework giving a generic syntax and the building blocks for defining specific restrictions thereof, called **theories**. The restrictions defined by a theory are on the nonlogical symbols that can be used and the interpretation that we can give them. Indeed, in a practical setting we would not want to consider an arbitrary interpretation of the symbols as above (where “+” is multiplication); rather we consider only specific ones.

A set of nonlogical symbols is called a **signature**. Given a signature  $\Sigma$ , a  $\Sigma$ -**formula** is a formula that uses only nonlogical symbols from  $\Sigma$  (possibly in addition to logical symbols). A variable is **free** if it is not bound by a quantifier. A **sentence** is a formula without free variables. A first-order  $\Sigma$ -**theory**  $T$  consists of a set of  $\Sigma$ -sentences. For a given  $\Sigma$ -theory  $T$ , a  $\Sigma$ -formula  $\varphi$  is  $T$ -**satisfiable** if there exists a structure that satisfies both the formula and the sentences of  $T$ . Similarly, a  $\Sigma$ -formula  $\varphi$  is  $T$ -**valid** if all structures that satisfy the sentences of  $T$ , also satisfy  $\varphi$ .

The set of sentences that are required is sometimes large or even infinite. It is therefore common to define theories via a set of axioms, which implicitly represent all the sentences that can be inferred from them, using some sound and complete inference system for the logical symbols.

**Example 1.23.** Consider a simple signature  $\Sigma$  consisting only of the predicate symbol “=”.<sup>6</sup> Let  $T$  be a  $\Sigma$ -theory. An example of a well-formed  $\Sigma$ -formula is

$$\forall x, y, z. ((x = y) \wedge \neg(y = z)) \implies \neg(x = z). \quad (1.34)$$

If we wish  $T$  to restrict the interpretation of “=” to the equality predicate, the following three axioms are sufficient:

$$\begin{aligned} \forall x. x = x & \quad (\text{REFLEXIVITY}), \\ \forall x, y. x = y \implies y = x & \quad (\text{SYMMETRY}), \\ \forall x, y, z. x = y \wedge y = z \implies x = z & \quad (\text{TRANSITIVITY}). \end{aligned} \quad (1.35)$$

Since every domain and interpretation that satisfy these axioms also satisfy (1.34), then (1.34) is  $T$ -valid.  $\blacksquare$

As said above, a theory restricts only the nonlogical symbols. If we want to restrict the set of logical symbols or the grammar, we need to speak about **fragments** of the logic. For example, we can speak about the **quantifier-free fragment** of  $T$  as defined in the example above. This fragment, called **equality logic**, happens to be the subject of Chap. 4. Most of the chapters, in fact, are concerned with quantifier-free fragments of theories. Another useful fragment is called the **conjunctive fragment**, which means that the only Boolean connective that is allowed is conjunction. What about restricting the interpretation of the logical symbols? The axioms that restrict the interpretation of the logical symbols, called the **logical axioms**, are assumed to be “built in”, i.e., they are common to all first-order theories.

Numerous theories have been considered over the years, corresponding to various problems of interest. Many of them lead to decidability, and, frequently to efficient decision procedures. The theory of **Presburger arithmetic**, for example, is defined with a signature  $\Sigma = \{0, 1, +, =\}$  and is still decidable. In

<sup>6</sup> It is frequently the case in the literature that the equality sign is considered as a logical symbol, and then the theory defined here has an empty signature. We do not follow this convention here, however.

contrast, the theory of **Peano arithmetic**, which is defined over a signature  $\Sigma = \{0, 1, +, \cdot, =\}$ , is undecidable. Thus, the addition of the multiplication symbol and the corresponding axioms that define it make the decision problem undecidable. Other famous theories include the theory of equality, the theory of reals, the theory of integers, the theory of arrays, the theory of recursive data structures and the theory of sets (“set theory”). Many of the decidable ones that are in practical use are covered in this book.

### 1.4.1 The Problem We Solve

Unless otherwise stated, we are concerned with

*the satisfiability problem of the quantifier-free fragment of various first-order theories.*

Formulas in such fragments are called **ground formulas**, as they only contain free (unquantified, also called ground) variables and constants. Exceptions are Chap. 9, which is concerned with quantified formulas, and a small part of Chap. 7, which is concerned with quantified array logic.

There is a subtle difference between the satisfiability problem of ground formulas and the satisfiability problem of existentially quantified formulas. It is, of course, trivial that a ground formula  $\varphi$  over variables  $x_1, \dots, x_n$  is satisfiable if and only if

$$\exists x_1, \dots, x_n. \varphi \tag{1.36}$$

is satisfiable. Thus, the decision procedures for both problems can be similar. The reason we use the former definition is that this entails, from a formal perspective, that the satisfying structure includes an assignment of the variables, because they are all free. In many practical applications, such an assignment is necessary. In fact, the former problem can be seen as an instance of the **constraint satisfaction problem (CSP)**, which is all about finding an assignment that satisfies a set of unquantified constraints.<sup>7</sup>

We assume that the input formulas are given in negation normal form, or that they are implicitly transformed to this form as a first step of any of the algorithms described later. As explained in Sect. 1.10, every formula can be transformed to this form in linear time. The reason that this assumption is important is that it simplifies the algorithms and the arguments for their correctness.

### 1.4.2 Our Presentation of Theories

Our presentation of theories in the chapters to come is not as defined above. In an attempt to make the presentation more accessible and the chapters more self-contained, we make the following changes:

<sup>7</sup> The emphasis and terminology are somewhat different. Most of the research in the CSP community is concerned with finite, discrete domains, in contrast to the problems considered in this book.



1. Rather than specifying theories through their set of symbols and sentences, we give the domain explicitly, and fix the interpretations of symbols in accordance with their common use. Hence, “+” is always the addition function, “0” is the 0 element in the given domain, and so forth.
2. Rather than specifying the theory fragment we are concerned with by referring to the generic grammar of first-order logic as a starting point, we give an explicit, self-contained definition of the grammar.

From a formal-logic point of view, fixing the interpretation means only that we have the sentences implicitly; the satisfiability problem remains the same. From the *algorithmic* point of view, however, the satisfiability problem now amounts to searching for a satisfying assignment of variables from the predefined domain. Whether a given assignment satisfies the formula can be determined according to the commonly used meanings of the various symbols.

This form of presentation is in line with our focus on the algorithmic point of view: when designing a decision procedure for a theory, the interpretation of the symbols has to be predefined. In other words, changing the domain or interpretation of symbols changes the algorithm.

## 1.5 Expressiveness vs. Decidability

There is an important trade-off between what a theory can express and how hard it is to decide, that is, how hard it is to determine whether a given formula allowed by the theory is valid or not. This is the reason for defining many different theories: otherwise, we would define and use only a single theory sufficiently expressive for all perceivable decision problems.

A theory can be seen as a tool for defining **languages**. Every formula in the theory defines a language, which is the set of “words” (the assignments, in the case of quantifier-free formulas) that satisfy it. We now define what it means that one theory is more expressive than another.

**Definition 1.24 (expressiveness).** *Theory A is more expressive than theory B if every language that can be defined by a B-formula can also be defined by an A-formula, and there exists at least one language definable by an A-formula that cannot be defined by a B-formula. We denote the fact that theory B is less expressive than theory A by  $B \prec A$ .*

$B \prec A$

For example, propositional logic is more expressive than what is known as “2-CNF”, i.e., CNF in which each clause has at most two literals. In propositional logic, we can define the formula

$$x_1 \vee x_2 \vee x_3 , \tag{1.37}$$

which defines a language that we cannot define with 2-CNF: it accepts all truth assignments to  $x_1, x_2, x_3$  except  $\{x_1 \mapsto \text{FALSE}, x_2 \mapsto \text{FALSE}, x_3 \mapsto \text{FALSE}\}$ . How can we prove this?

Well, assume that there exists a 2-CNF representation of this formula using the same set of variables, and consider one of its binary clauses. Such a clause contradicts two of the eight possible assignments. For example, a clause  $(x_1 \vee x_2)$  contradicts  $\{x_1 \mapsto \text{FALSE}, x_2 \mapsto \text{FALSE}, x_3 \mapsto \text{FALSE}\}$  and  $\{x_1 \mapsto \text{FALSE}, x_2 \mapsto \text{FALSE}, x_3 \mapsto \text{TRUE}\}$ . Any additional clause can only contradict more assignments. Hence, we can never create a 2-CNF formula such that exactly one of the eight assignments does not satisfy it.

On the other hand, 2-CNF is a restriction of propositional logic; hence, obviously, any 2-CNF formula can be expressed in propositional logic. Thus, we have

$$\text{2-CNF} \prec \text{propositional logic} . \quad (1.38)$$

This example also demonstrates the influence of expressiveness on computational hardness: while propositional logic is NP-complete, 2-CNF can be solved in polynomial time.


In order to illustrate the trade-off between how expressive a theory is and how hard it is to decide formulas in that theory, consider a theory  $T$  defined by some syntax. Let  $T_1, \dots, T_n$  denote a list of fragments of  $T$ , defined by various restrictions on the syntax of  $T$  (similarly to the way we restricted propositional logic to 2-CNF above), for which  $T_1 \prec T_2 \prec \dots \prec T_n \prec T$ . Technically, this means that we have imposed a **total order** on these fragments in terms of their expressive power. Under such assumptions, Fig. 1.3 illustrates the trade-off between expressiveness and computational hardness: the less expressive the theory is (the more restrictions we put on it), the easier it is to decide it. Assume our imaginary theory  $T$  is undecidable. After some threshold is crossed (from right to left in the figure), the theory fragments can become decidable. After enough restrictions have been added, the theory becomes solvable in polynomial time. The decidable but nonpolynomially decidable fragments pose a *computational challenge*. This is one of the challenges we focus on in this book.

This view is simplistic, however, because there is no total order on the expressive power of theories, only a partial order. This means that there can be two theories,  $A$  and  $B$ , neither of which is more expressive than the other, yet their expressive power is different. In other words, there are languages that can be defined by  $A$  and not by  $B$ , and there are languages that can be defined by  $B$  and not by  $A$ .

## 1.6 Boolean Structure in Decision Problems

We conclude this chapter by demonstrating the need for reasoning about formulas with a Boolean structure.

Many decision procedures assume that the decision problem is given by a conjunction of constraints. The simplex algorithm and the Omega test, both of which are described in Chap. 5, are examples of such procedures.



**Fig. 1.3.** The trade-off between expressiveness of theories and the hardness of deciding them, illustrated for an imaginary series of theories  $T_1, \dots, T_n, T$  for which each  $T_i$ ,  $i \in \{1, \dots, n\}$ , is less expressive than its successor

Many applications, however, require a more complex Boolean structure. In program analysis and verification, for example, disjunctions may appear in the program to be verified, either explicitly (e.g.,  $x = y \mid\mid z$ ) or implicitly through constructs such as `if` and `switch` statements. Any reasoning system about such programs, therefore, must be able to deal with disjunctions. For example, **verification conditions** that arise in program verification (e.g., using **Hoare logic**), often have the form of an implication.

The following example focuses on a technique for reasoning about programs, that demonstrates how program structure, including `if` statements, is evident in the underlying verification conditions that need to be checked.

**Example 1.25. Bounded model checking (BMC)** of programs is a technique for verifying that a given property (typically given as an assertion by the user) holds for a program in which the number of loop iterations and recursive calls is bounded by a given number  $k$ . The states that the program can reach within this bound are represented symbolically by a formula, together with the negation of the property. If the combined formula is satisfiable, then there exists a path in the program that violates the property.

Consider the program in the left part of Fig. 1.4. The number of paths through this program is exponential in  $N$ , as each of the  $a[i]$  elements can be either zero or nonzero. Despite the exponential number of paths through the program, its states can be encoded with a formula of size linear in  $N$ , as demonstrated in the right part of the figure.

The formula on the right of Fig. 1.4 encodes the states of the program on its left, using the **static-single-assignment (SSA)** form. Briefly, this means that in each assignment of the form  $x = \text{exp}$ ;, the left-hand side variable  $x$  is replaced with a new variable, say  $x_1$ , and any reference to  $x$  after this line and before  $x$  is assigned again is replaced with  $x_1$ . Such a replacement is possible because there are no loops (recall that this is done in the context of BMC). After this transformation, the statements are conjoined. The resulting equation represents the states of the original program.

<pre> int a[N]; unsigned c; ... c = 0; for(i = 0; i &lt; N; i++)     if(a[i] == 0)         c++; </pre>	$ \begin{aligned} c_1 &= 0 \wedge \\ c_2 &= (a[0] = 0) ? c_1 + 1 : c_1 \wedge \\ c_3 &= (a[1] = 0) ? c_2 + 1 : c_2 \wedge \\ &\dots \\ c_{N+1} &= (a[N-1] = 0) ? c_N + 1 : c_N \end{aligned} $
--	--

**Fig. 1.4.** A simple program with an exponential number of paths (*left*), and a static-single-assignment (SSA) form of this program after unwinding its `for` loop (*right*)

The ternary operator  $c ? x : y$  in the equation on the right of Fig. 1.4 can be rewritten using a disjunction, as illustrated in (1.39). These disjunctions lead to an exponential number of clauses once the formula is converted to DNF.

$$\begin{aligned}
&c_1 = 0 \wedge \\
&((a[0] = 0 \wedge c_2 = c_1 + 1) \vee (a[0] \neq 0 \wedge c_2 = c_1)) \wedge \\
&((a[1] = 0 \wedge c_3 = c_2 + 1) \vee (a[1] \neq 0 \wedge c_3 = c_2)) \wedge \\
&\dots \\
&((a[N-1] = 0 \wedge c_{N+1} = c_N + 1) \vee (a[N-1] \neq 0 \wedge c_{N+1} = c_N)) .
\end{aligned} \tag{1.39}$$

In order to verify that some assertion holds at a specific location in the program, it is sufficient to add a constraint corresponding to the negation of this assertion, and check whether the resulting formula is satisfiable. For example, to prove that at the end of the program  $c \leq N$ , we need to conjoin (1.39) with  $(c_{N+1} > N)$ . ▀

To summarize this section, there is a need to reason about formulas with disjunctions, as illustrated in the example above. The simple solution of going through DNF does not scale, and better solutions are needed. Solutions that perform better in practice (the worst case remains exponential, of course) indeed exist, and are covered extensively in this book.

## 1.7 Problems

### Problem 1.1 (improving Tseitin's encoding).

- (a) Using Tseitin's encoding, transform the following formula  $\varphi$  to CNF. How many clauses are needed?

$$\varphi := \neg(x_1 \wedge (x_2 \vee \dots \vee x_n)) . \tag{1.40}$$

- (b) Consider a clause  $(x_1 \vee \dots \vee x_n)$ ,  $n > 2$ , in a non-CNF formula. How many auxiliary variables are necessary for encoding it with Tseitin's encoding? Suggest an alternative way to encode it, using a single auxiliary variable. How many clauses are needed?

**Problem 1.2 (expressiveness and complexity).**

- (a) Let  $T_1$  and  $T_2$  be two theories whose satisfiability problem is decidable and in the same complexity class. Is the satisfiability problem of a  $T_1$ -formula reducible to a satisfiability problem of a  $T_2$ -formula?
- (b) Let  $T_1$  and  $T_2$  be two theories whose satisfiability problems are reducible to one another. Are  $T_1$  and  $T_2$  in the same complexity class?

**Problem 1.3 (monotonicity of NNF with respect to satisfiability).** Prove Theorem 1.14.

**Problem 1.4 (one-sided Tseitin encoding).** Let  $\varphi$  be an NNF formula (see Definition 1.10). Let  $\vec{\varphi}$  be a formula derived from  $\varphi$  as in Tseitin's encoding (see Sect. 1.3), but where the CNF constraints are derived from implications from left to right rather than equivalences. For example, given a formula

$$a_1 \wedge (a_2 \vee \neg a_3),$$

the new encoding is the CNF equivalent of the following formula,

$$\begin{aligned} & x_0 && \wedge \\ (x_0 & \implies a_1 \wedge x_1) && \wedge \\ (x_1 & \implies a_2 \vee x_2) && \wedge \\ (x_2 & \implies \neg a_3), \end{aligned}$$

where  $x_0, x_1, x_2$  are new auxiliary variables. Note that Tseitin's encoding to CNF starts with the same formula, except that the " $\implies$ " symbol is replaced with " $\iff$ ".

1. Prove that  $\vec{\varphi}$  is satisfiable if and only if  $\varphi$  is.
2. Let  $l, m, n$  be the number of AND, OR, and NOT gates, respectively, in  $\varphi$ . Derive a formula parameterized by  $l, m$  and  $n$  that expresses the ratio of the number of CNF clauses in Tseitin's encoding to that in the one-sided encoding suggested here.

## 1.8 Glossary

The following symbols were used in this chapter:

<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$\alpha \models \varphi$	An assignment $\alpha$ satisfies a formula $\varphi$	6
$\models \varphi$	A formula $\varphi$ is valid (in the case of quantifier-free formulas, this means that it is satisfied by all assignments from the domain)	6
$T$	A theory	6
$pos(\alpha, \varphi)$	Set of literals of $\varphi$ satisfied by an assignment $\alpha$	9
$B \prec A$	Theory $B$ is less expressive than theory $A$	18

---

## Decision Procedures for Propositional Logic

### 2.1 Propositional Logic

We assume that the reader is familiar with propositional logic. The syntax of formulas in propositional logic is defined by the following grammar:

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid \neg \text{formula} \mid (\text{formula}) \mid \text{atom} \\ \text{atom} &: \text{Boolean-identifier} \mid \text{TRUE} \mid \text{FALSE} \end{aligned}$$

Other Boolean operators such as OR ( $\vee$ ) can be constructed using AND ( $\wedge$ ) and NOT ( $\neg$ ).

#### 2.1.1 Motivation

Propositional logic is widely used in diverse areas such as database queries, planning problems in artificial intelligence, automated reasoning and circuit design. Here we consider two examples: a layout problem and a program verification problem.

**Example 2.1.** Let  $S = \{s_1, \dots, s_n\}$  be a set of radio stations, each of which has to be allocated one of  $k$  transmission frequencies, for some  $k < n$ . Two stations that are too close to each other cannot have the same frequency. The set of pairs having this constraint is denoted by  $E$ . To model this problem, define a set of propositional variables  $\{x_{ij} \mid i \in \{1, \dots, n\}, j \in \{1, \dots, k\}\}$ . Intuitively, variable  $x_{ij}$  is set to TRUE if and only if station  $i$  is assigned the frequency  $j$ . The constraints are:

- Every station is assigned at least one frequency:

$$\bigwedge_{i=1}^n \bigvee_{j=1}^k x_{ij} . \tag{2.1}$$

- Every station is assigned not more than one frequency:

$$\bigwedge_{i=1}^n \bigwedge_{j=1}^{k-1} (x_{ij} \implies \bigwedge_{j < t \leq k} \neg x_{it}). \quad (2.2)$$

- Close stations are not assigned the same frequency. For each  $(i, j) \in E$ ,

$$\bigwedge_{t=1}^k (x_{it} \implies \neg x_{jt}). \quad (2.3)$$

Note that the input of this problem can be represented by a graph, where the stations are the graph's nodes and  $E$  corresponds to the graph's edges. Checking whether the allocation problem is solvable corresponds to solving what is known in graph theory as the *k-colorability* problem: can all nodes be assigned one of  $k$  colors such that two adjacent nodes are assigned different colors? Indeed, one way to solve *k-colorability* is by reducing it to propositional logic. ▀

**Example 2.2.** Consider the two code fragments in Fig. 2.1. The fragment on the right-hand side might have been generated from the fragment on the left-hand side by an optimizing compiler.

<pre>if(!a &amp;&amp; !b) h(); else     if(!a) g();     else f();</pre>	<pre>if(a) f(); else     if(b) g();     else h();</pre>
---	---

**Fig. 2.1.** Two code fragments – are they equivalent?

We would like to check if the two programs are equivalent. The first step in building the **verification condition** is to model the variables  $a$  and  $b$  and the procedures that are called using the Boolean variables  $a$ ,  $b$ ,  $f$ ,  $g$ , and  $h$ , as can be seen in Fig. 2.2.

<pre>if ¬a ∧ ¬b then h else     if ¬a then g     else f</pre>	<pre>if a then f else     if b then g     else h</pre>
---	--

**Fig. 2.2.** In the process of building a formula – the verification condition – we replace the program variables and the function symbols with new Boolean variables

The *if-then-else* construct can be replaced by an equivalent propositional logic expression as follows:



$$(\text{if } x \text{ then } y \text{ else } z) \equiv (x \wedge y) \vee (\neg x \wedge z). \quad (2.4)$$

Consequently, the problem of checking the equivalence of the two code fragments is reduced to checking the validity of the following propositional formula:

$$\iff (\neg a \wedge \neg b) \wedge h \vee \neg(\neg a \wedge \neg b) \wedge (\neg a \wedge g \vee a \wedge f) \quad (2.5)$$

$$\iff a \wedge f \vee \neg a \wedge (b \wedge g \vee \neg b \wedge h).$$

▀

## 2.2 SAT Solvers

### 2.2.1 The Progress of SAT Solving


Given a Boolean formula  $\mathcal{B}$ , a SAT solver decides whether  $\mathcal{B}$  is satisfiable; if it is, it also reports a satisfying assignment. In this chapter, we consider only the problem of solving formulas in conjunctive normal form (CNF) (see Definition 1.20). Since every formula can be converted to this form in linear time (as explained right after Definition 1.20), this does not impose a real restriction.<sup>1</sup> Solving general propositional formulas can be somewhat more efficient in some problem domains, but most of the solvers and most of the research are still focused on CNF formulas.

The practical and theoretical importance of the satisfiability problem has led to a vast amount of research in this area, which has resulted in exceptionally powerful SAT solvers. Modern SAT solvers can solve many real-life CNF formulas with hundreds of thousands or even millions of variables in a reasonable amount of time. Figure 2.3 shows a sketch of the progress of these tools through the years. Of course, there are also instances of problems two orders of magnitude smaller that these tools still cannot solve. In general, it is very hard to predict which instance is going to be hard to solve, without actually attempting to solve it.

For many years, SAT solvers were better at solving satisfiable instances than unsatisfiable ones. This is not true anymore. The success of SAT solvers can be largely attributed to their ability to learn from wrong assignments, to prune large search spaces quickly, and to focus first on the “important” variables, those variables that, once given the right value, simplify the problem immensely.<sup>2</sup> All of these factors contribute to the fast solving of both satisfiable and unsatisfiable instances.

<sup>1</sup> Appendix B provides a library for performing this conversion and generating CNF in the DIMACS format, which is used by virtually all publicly available SAT solvers.

<sup>2</sup> Specifically, every formula has what is known as **back-door variables** [200], which are variables that, once given the right value, simplify the formula to the point that it is polynomial to solve.



**Fig. 2.3.** The size of industrial CNF formulas (instances generated for solving various realistic problems such as verification of circuits and planning problems) that are regularly solved by SAT solvers in a few hours, according to year. Most of the progress in efficiency has been made in the last decade

The majority of modern SAT solvers can be classified into two main categories. The first category is based on the *Davis–Putnam–Loveland–Logemann* (**DPLL**) framework: in this framework the tool can be thought of as traversing and backtracking on a binary tree, in which internal nodes represent partial assignments, and the leaves represent full assignments, i.e., an assignment to all the variables.

The second category is based on a **stochastic search**: the solver guesses a full assignment, and then, if the formula is evaluated to FALSE under this assignment, starts to flip values of variables according to some (greedy) heuristic. Typically it counts the number of unsatisfied clauses and chooses the flip that minimizes this number. There are various strategies that help such solvers avoid local minima and avoid repeating previous bad moves. DPLL solvers, however, are considered better in most cases, at least at the time of writing this chapter (2007), according to annual competitions that measure their performance with numerous CNF instances. DPLL solvers also have the advantage that, unlike most stochastic search methods, they are complete (see Definition 1.6). Stochastic methods seem to have an average advantage in solving randomly generated (satisfiable) CNF instances, which is not surprising: in these instances there is no structure to exploit and learn from, and no obvious choices of variables and values, which makes the heuristics adopted by DPLL solvers ineffective. We shall focus on DPLL solvers only.

### 2.2.2 The DPLL Framework

In its simplest form, a DPLL solver progresses by making a decision about a variable and its value, propagates implications of this decision that are easy to detect, and backtracks in the case of a conflict. Viewing the process as a

search on a binary tree, each decision is associated with a **decision level**, which is the depth in the binary decision tree in which it is made, starting from 1. The assignments implied by a decision are associated with its decision level. Assignments implied regardless of the current assignments (owing to **unary clauses**, which are clauses with a single literal) are associated with decision level 0, also called the **ground level**.

**Definition 2.3 (state of a clause under an assignment).** *A clause is **satisfied** if one or more of its literals are satisfied (see Definition 1.12), **conflicting** if all of its literals are assigned but not satisfied, **unit** if it is not satisfied and all but one of its literals are assigned, and **unresolved** otherwise.*

Note that the definition of a unit clause and an unresolved clause are only relevant for partial assignments (see Definition 1.1).

**Example 2.4.** Given the partial assignment

$$\{x_1 \mapsto 1, x_2 \mapsto 0, x_4 \mapsto 1\}, \quad (2.6)$$

$(x_1 \vee x_3 \vee \neg x_4)$	is satisfied,
$(\neg x_1 \vee x_2)$	is conflicting,
$(\neg x_1 \vee \neg x_4 \vee x_3)$	is unit,
$(\neg x_1 \vee x_3 \vee x_5)$	is unresolved.

▀

Given a partial assignment under which a clause becomes unit, it must be extended so that it satisfies the unassigned literal of this clause. This observation is known as the **unit clause rule**. Following this requirement is necessary but obviously not sufficient for satisfying the formula.

For a given unit clause  $C$  with an unassigned literal  $l$ , we say that  $l$  is implied by  $C$  and that  $C$  is the **antecedent clause** of  $l$ , denoted by  $Antecedent(l)$ . If more than one unit clause implies  $l$ , we refer to the clause that the SAT solver used in order to imply  $l$  as its antecedent.

**Example 2.5.** The clause  $C := (\neg x_1 \vee \neg x_4 \vee x_3)$  and the partial assignment  $\{x_1 \mapsto 1, x_4 \mapsto 1\}$ , imply the assignment  $x_3$  and  $Antecedent(x_3) = C$ . ▀

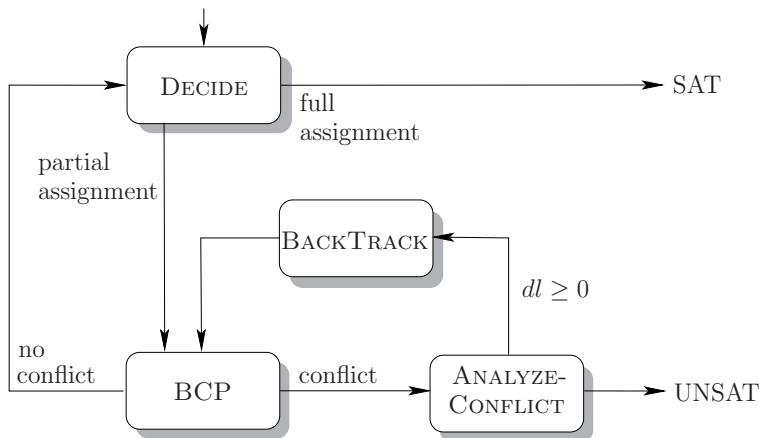
A framework followed by most modern DPLL solvers has been presented by, for example, Zhang and Malik [211], and is shown in Algorithm 2.2.1. The table in Fig. 2.5 includes a description of the main components used in this algorithm, and Fig. 2.4 depicts the interaction between them. A description of the ANALYZE-CONFLICT function is delayed to Sect. 2.2.6.

**Algorithm 2.2.1:** DPLL-SAT**Input:** A propositional CNF formula  $\mathcal{B}$ **Output:** “Satisfiable” if the formula is satisfiable and “Unsatisfiable” otherwise

```

1. function DPLL
2.   if BCP() = “conflict” then return “Unsatisfiable”;
3.   while (TRUE) do
4.     if  $\neg$ DECIDE() then return “Satisfiable”;
5.     else
6.       while (BCP() = “conflict”) do
7.          $backtrack\text{-}level :=$  ANALYZE-CONFLICT();
8.         if  $backtrack\text{-}level < 0$  then return “Unsatisfiable”;
9.         else BackTrack( $backtrack\text{-}level$ );

```



**Fig. 2.4.** DPLL-SAT: high-level overview of the Davis-Putnam-Loveland-Logemann algorithm. The variable  $dl$  is the decision level to which the procedure backtracks

### 2.2.3 BCP and the Implication Graph

We now demonstrate Boolean constraints propagation (BCP), reaching a conflict, and backtracking. Each assignment is associated with the decision level at which it occurred. If a variable  $x_i$  is assigned 1 (TRUE) (owing to either a decision or an implication) at decision level  $dl$ , we write  $x_i@dl$ . Similarly,  $\neg x_i@dl$  reflects an assignment of 0 (FALSE) to this variable at decision level  $dl$ . Where appropriate, we refer only to the truth assignment, omitting the decision level, in order to make the notation simpler.

$x_i@dl$

<b>Name</b>	DECIDE()
<i>Output</i>	FALSE if and only if there are no more variables to assign.
<i>Description</i>	Chooses an unassigned variable and a truth value for it.
<i>Comments</i>	There are numerous heuristics for making these decisions, some of which are described later in Sect. 2.2.5. Each such decision is associated with a decision level, which can be thought of as the depth in the search tree.
<b>Name</b>	BCP()
<i>Output</i>	“conflict” if and only if a conflict is encountered.
<i>Description</i>	Repeated application of the unit clause rule until either a conflict is encountered or there are no more implications.
<i>Comments</i>	This repeated process is called Boolean constraint propagation (BCP). BCP is applied in line 2 because unary clauses at this stage are unit clauses.
<b>Name</b>	ANALYZE-CONFLICT()
<i>Output</i>	Minus 1 if a conflict at decision level 0 is detected (which implies that the formula is unsatisfiable). Otherwise, a decision level which the solver should backtrack to.
<i>Description</i>	A detailed description of this function is delayed to Sect. 2.2.4. Briefly, it is responsible for computing the backtracking level, detecting global unsatisfiability, and adding new constraints on the search in the form of new clauses.
<b>Name</b>	BACKTRACK( <i>dl</i> )
<i>Description</i>	Sets the current decision level to <i>dl</i> and erases assignments at decision levels larger than <i>dl</i> .

**Fig. 2.5.** A description of the main components of Algorithm 2.2.1

The process of BCP is best illustrated with an **implication graph**. An implication graph represents the current partial assignment and the reason for each of the implications.

**Definition 2.6 (implication graph).** *An implication graph is a labeled directed acyclic graph  $G(V, E)$ , where:*

- *$V$  represents the literals of the current partial assignment (we refer to a node and the literal that it represents interchangeably). Each node is labeled with the literal that it represents and the decision level at which it entered the partial assignment.*
- *$E$  with  $E = \{(v_i, v_j) \mid v_i, v_j \in V, \neg v_i \in \text{Antecedent}(v_j)\}$  denotes the set of directed edges where each edge  $(v_i, v_j)$  is labeled with  $\text{Antecedent}(v_j)$ .*

- $G$  can also contain a single **conflict node** labeled with  $\kappa$  and incoming edges  $\{(v, \kappa) \mid \neg v \in c\}$  labeled with  $c$  for some conflicting clause  $c$ .

The root nodes of an implication graph correspond to decisions, and the internal nodes to implications through BCP. A conflict node with incoming edges labeled with  $c$  represents the fact that the BCP process has reached a conflict, by assigning 0 to all the literals in the clause  $c$  (i.e.,  $c$  is conflicting). In such a case, we say that the graph is a **conflict graph**. The implication graph corresponds to all the decision levels lower than or equal to the current one, and is dynamic: backtracking removes nodes and their incoming edges, whereas new decisions, implications, and conflict clauses increase the size of the graph.

The implication graph is sensitive to the order in which the implications are propagated in BCP, which means that the graph is not unique for a given partial assignment. In most SAT solvers, this order is rather arbitrary (in particular, BCP progresses along a list of clauses that contain a given literal, and the order of clauses in this list can be sensitive to the order of clauses in the input CNF formula). In some other SAT solvers – see for example [151] – this order is not arbitrary; rather, it is biased towards reaching a conflict faster.


A **partial implication graph** is a subgraph of an implication graph, which illustrates the BCP at a specific decision level. Partial implication graphs are sufficient for describing ANALYZE-CONFLICT. The roots in such a partial graph represent assignments (not necessarily decisions) at decision levels lower than  $dl$ , in addition to the decision at level  $dl$ , and internal nodes correspond to implications at level  $dl$ . The description that follows uses mainly this restricted version of the graph.

Consider, for example, a formula that contains the following set of clauses, among others:

$$\begin{aligned}
 c_1 &= (\neg x_1 \vee x_2), \\
 c_2 &= (\neg x_1 \vee x_3 \vee x_5), \\
 c_3 &= (\neg x_2 \vee x_4), \\
 c_4 &= (\neg x_3 \vee \neg x_4), \\
 c_5 &= (x_1 \vee x_5 \vee \neg x_2), \\
 c_6 &= (x_2 \vee x_3), \\
 c_7 &= (x_2 \vee \neg x_3), \\
 c_8 &= (x_6 \vee \neg x_5).
 \end{aligned} \tag{2.7}$$

Assume that at decision level 3 the decision was  $\neg x_6@3$ , which implied  $\neg x_5@3$  owing to clause  $c_8$  (hence,  $Antecedent(\neg x_5) = c_8$ ). Assume further that the solver is now at decision level 6 and assigns  $x_1 = 1$ . At decision levels 4 and 5, variables other than  $x_1, \dots, x_6$  were assigned, and are not listed here as they are not relevant to these clauses.

The implication graph on the left of Fig. 2.6 demonstrates the BCP process at the current decision level 6 until, in this case, a conflict is detected. The roots of this graph, namely  $\neg x_5@3$  and  $x_1@6$ , constitute a sufficient condition



**Fig. 2.6.** A partial implication graph for decision level 6, corresponding to the clauses in (2.7), after a decision  $x_1 = 1$  (*left*) and a similar graph after learning the conflict clause  $c_9 = (x_5 \vee \neg x_1)$  and backtracking to decision level 3 (*right*)

for creating this conflict. Therefore, we can safely add to our formula the **conflict clause**

$$c_9 = (x_5 \vee \neg x_1). \quad (2.8)$$

While  $c_9$  is logically implied by the original formula and therefore does not change the result, it prunes the search space. The process of adding conflict clauses is generally referred to as **learning**, reflecting the fact that this is the solver's way to learn from its past mistakes. As we progress in this chapter, it will become clear that conflict clauses not only prune the search space, but also have an impact on the decision heuristic, the backtracking level, and the set of variables implied by each decision.

ANALYZE-CONFLICT is the function responsible for deriving new conflict clauses and computing the backtracking level. It traverses the implication graph backwards, starting from the conflict node  $\kappa$ , and generates a conflict clause through a series of steps that we describe later in Sect. 2.2.4. For now, assume that  $c_9$  is indeed the clause generated.

After detecting the conflict and adding  $c_9$ , the solver determines which decision level to backtrack to according to the **conflict-driven backtracking** strategy. According to this strategy, the backtracking level is set to the *second most recent decision level in the conflict clause*<sup>3</sup> (or, equivalently, it is set to the highest of the decision levels in the clause other than the current decision level), while erasing all decisions and implications made *after* that level.

In the case of  $c_9$ , the solver backtracks to decision level 3 (the decision level of  $x_5$ ), and erases all assignments from decision level 4 onwards, including the assignments to  $x_1, x_2, x_3$ , and  $x_4$ .

The newly added conflict clause  $c_9$  becomes a unit clause since  $x_5 = 0$ , and therefore the assignment  $\neg x_1@3$  is implied. This new implication re-starts the BCP process at level 3. Clause  $c_9$  is a special kind of a conflict clause, called an **asserting clause**: it forces an immediate implication after backtracking. ANALYZE-CONFLICT can be designed to generate asserting clauses only, as indeed most competitive solvers do.

<sup>3</sup> In the case of learning a unary clause, the solver backtracks to the ground level.

**Aside: Multiple Conflict Clauses**

More than one conflict clause can be derived from a conflict graph. In the present example, the assignment  $\{x_2 \mapsto 1, x_3 \mapsto 1\}$  is also a sufficient condition for the conflict, and hence  $(\neg x_2 \vee \neg x_3)$  is also a conflict clause. A generalization of this observation requires the following definition.

**Definition 2.7 (separating cut).** *A separating cut in a conflict graph is a minimal set of edges whose removal breaks all paths from the root nodes to the conflict node.*

This definition is applicable to a full implication graph (see Definition 2.6), as well as to a partial graph focused on the decision level of the conflict. The cut bipartitions the nodes into the *reason* side (the side that includes all the roots) and the *conflict* side. The set of nodes on the reason side that have at least one edge to a node on the conflict side constitute a sufficient condition for the conflict, and hence their negation is a legitimate conflict clause. Different SAT solvers have different strategies for choosing the conflict clauses that they add: some add as many as possible (corresponding to many different cuts), while others try to find the most effective ones. Some, including most of the modern SAT solvers, add a single clause, which is an asserting clause (see below), for each conflict.

After asserting  $x_1 = 0$  the solver again reaches a conflict, as can be seen in the right drawing in Fig. 2.6. This time the conflict clause  $(x_2)$  is added, the solver backtracks to decision level 0, and continues from there. Why  $(x_2)$ ? The strategy of ANALYZE-CONFLICT in generating these clauses is explained later in Sect. 2.2.4, but observe for the moment how indeed  $\neg x_2$  leads to a conflict through clauses  $c_6$  and  $c_7$ , as can also be inferred from Fig. 2.6 (right).

Conflict-driven backtracking raises several issues:

- *It seems to waste work*, because the partial assignments up to decision level 5 can still be part of a satisfying assignment. However, empirical evidence shows that conflict-driven backtracking, coupled with a conflict-driven decision heuristic such as VSIDS (discussed later in Sect. 2.2.5), performs very well. A possible explanation for the success of this heuristic is that the conflict encountered can influence the decision heuristic to decide values or variables different from those at deeper decision levels (levels 4 and 5 in this case). Thus, keeping the decisions and implications made before the new information (i.e., the new conflict clause) has arrived may skew the search to areas not considered best anymore by the heuristic. There has been some success in overcoming this problem by repeating previous assignments – see [150].
- *Is this process guaranteed to terminate?* In other words, how do we know that a partial assignment cannot be repeated forever? The learned conflict clauses cannot be the reason, because in fact most SAT solvers erase many

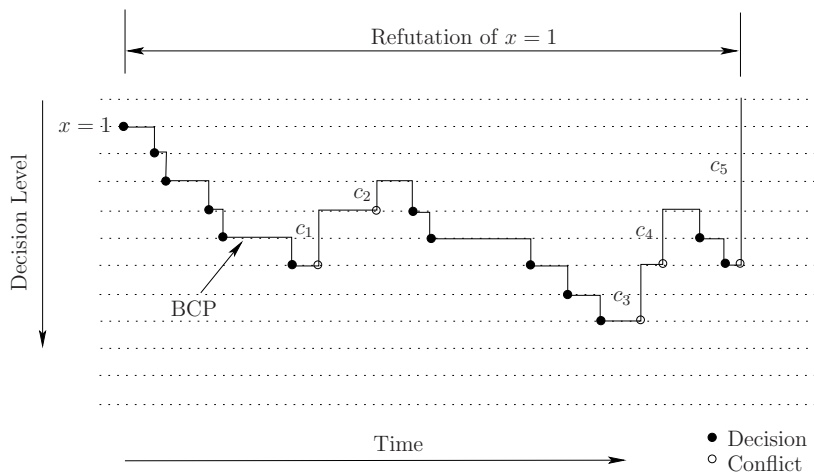


of them after a while to prevent the formula from growing too much. The reason is the following.

**Theorem 2.8.** *It is never the case that the solver enters decision level  $dl$  again with the same partial assignment.*

*Proof.* Consider a partial assignment up to decision level  $dl - 1$  that does not end with a conflict, and assume falsely that this state is repeated later, after the solver backtracks to some lower decision level  $dl^-$  ( $0 \leq dl^- < dl$ ). Any backtracking from a decision level  $dl^+$  ( $dl^+ \geq dl$ ) to decision level  $dl^-$  adds an implication at level  $dl^-$  of a variable that was assigned at decision level  $dl^+$ . Since this variable has not so far been part of the partial assignment up to decision level  $dl$ , once the solver reaches  $dl$  again, it is with a different partial assignment, which contradicts our assumption. ▀

The (hypothetical) progress of a SAT solver based on this strategy is illustrated in Fig. 2.7. More details of this graph are explained in the caption.



**Fig. 2.7.** Illustration of the progress of a SAT solver based on conflict-driven backtracking. Every conflict results in a conflict clause (denoted by  $c_1, \dots, c_5$  in the drawing). If the top left decision is  $x = 1$ , then this drawing illustrates the work done by the SAT solver to refute this wrong decision. Only some of the work during this time was necessary for creating  $c_5$ , refuting this decision, and computing the backtracking level. The “wasted work” (which might, after all, become useful later on) is due to the imperfection of the decision heuristic

### 2.2.4 Conflict Clauses and Resolution

Now consider ANALYZE-CONFLICT (Algorithm 2.2.2). The description of the algorithm so far has relied on the fact that the conflict clause generated is

an asserting clause, and we therefore continue with this assumption when considering the termination criterion for line 3. The following definitions are necessary for describing this criterion.

**Algorithm 2.2.2: ANALYZE-CONFLICT**

**Input:**

**Output:** Backtracking decision level + a new conflict clause

1. **if** *current-decision-level* = 0 **then return** -1;
2. *cl* := *current-conflicting-clause*;
3. **while** ( $\neg$ STOP-CRITERION-MET(*cl*)) **do**
4.     *lit* := LAST-ASSIGNED-LITERAL(*cl*);
5.     *var* := VARIABLE-OF-LITERAL(*lit*);
6.     *ante* := ANTECEDENT(*lit*);
7.     *cl* := RESOLVE(*cl*, *ante*, *var*);
8.     add-clause-to-database(*cl*);
9. **return** clause-asserting-level(*cl*);     ▷ 2nd highest decision level in *cl*

**Definition 2.9 (unique implication point (UIP)).** *Given a partial conflict graph corresponding to the decision level of the conflict, a unique implication point (UIP) is any node other than the conflict node that is on all paths from the decision node to the conflict node.*


The decision node itself is a UIP by definition, while other UIPs, if they exist, are internal nodes corresponding to implications at the decision level of the conflict.

**Definition 2.10 (first UIP).** *A first UIP is a UIP that is closest to the conflict node.*

We leave the proof that the notion of a first UIP in a conflict graph is well defined as an exercise (see Problem 2.11). Figure 2.8 demonstrates UIPs in a conflict graph (see also the caption).

Empirical studies show that a good strategy for the STOP-CRITERION-MET(*cl*) function (line 3) is to return TRUE if and only if *cl* contains the negation of the first UIP as its single literal at the current decision level. This negated literal becomes asserted immediately after backtracking. There are several advantages to this strategy, which may explain the results of the empirical studies:

1. *The strategy has a low computational cost, compared to strategies that choose UIPs further away from the conflict.*



**Fig. 2.8.** An implication graph (stripped of most of its labels) with two UIPs. The left UIP is the decision node, and the right one is the first UIP, as it is the one closest to the conflict node

## 2. It backtracks to the lowest decision level.

The second fact can be demonstrated with the help of Fig. 2.8. Let  $l_1$  and  $l_2$  denote the literals at the first and the second UIP, respectively. The asserting clauses generated with the first UIP and second-UIP strategies are, respectively,  $(\neg l_1 \vee \neg x_1 \vee \neg x_2)$  and  $(\neg l_2 \vee \neg x_1 \vee \neg x_2 \vee \neg x_4)$ . It is not a coincidence that the second clause subsumes the first, other than the asserting literals  $\neg l_1$  and  $\neg l_2$ : it is always like this, by construction. Now recall how the backtracking level is determined: it is equal to the decision level corresponding to the second highest in the asserting clause. Clearly, this implies that the backtracking level computed with regard to the first clause is lower than that computed with regard to the second clause. In our example, these are decision levels 4 and 7, respectively.

In order to explain lines 4–7 of ANALYZE-CONFLICT, we need the following definition.

**Definition 2.11 (binary resolution and related terms).** Consider the following inference rule:

$$\frac{(a_1 \vee \dots \vee a_n \vee \beta) \quad (b_1 \vee \dots \vee b_m \vee \neg\beta)}{(a_1 \vee \dots \vee a_n \vee b_1 \vee \dots \vee b_m)} \quad (\text{BINARY RESOLUTION}), \quad (2.9)$$

where  $a_1, \dots, a_n, b_1, \dots, b_m$  are literals and  $\beta$  is a variable. The variable  $\beta$  is called the **resolution variable**. The clauses  $(a_1 \vee \dots \vee a_n \vee \beta)$  and  $(b_1 \vee \dots \vee b_m \vee \neg\beta)$  are the **resolving clauses**, and  $(a_1 \vee \dots \vee a_n \vee b_1 \vee \dots \vee b_m)$  is the **resolvent clause**.

A well-known result obtained by Robinson [166] shows that a deductive system based on the binary-resolution rule as its single inference rule is sound and complete. In other words, a CNF formula is unsatisfiable if and only if there exists a finite series of binary-resolution steps ending with the empty clause.

**Aside: Hard Problems for Resolution-Based Procedures**


Some propositional formulas can be decided with no less than an exponential number of resolution steps in the size of the input. Haken [90] proved in 1985 that the **pigeonhole problem** is one such problem: given  $n > 1$  pigeons and  $n - 1$  pigeonholes, can each of the pigeons be assigned a pigeonhole without sharing? While a formulation of this problem in propositional logic is rather trivial with  $n \cdot (n - 1)$  variables, currently no SAT solver (which, recall, implicitly perform resolution) can solve this problem in a reasonable amount of time for  $n$  larger than several tens, although the size of the CNF itself is relatively small. As an experiment, we tried to solve this problem for  $n = 20$  with three leading SAT solvers: Siege4 [171], zChaff-04 [133] and HaifaSat [82]. On a Pentium 4 with 1 GB of main memory, none of the three could solve this problem within three hours. Compare this result with the fact that, bounded by the same timeout, these tools routinely solve problems arising in industry with hundreds of thousands of variables.

The function  $\text{RESOLVE}(c_1, c_2, v)$  used in line 7 of `ANALYZE-CONFLICT` returns the resolvent of the clauses  $c_1, c_2$ , where the resolution variable is  $v$ . The `ANTECEDENT` function used in line 6 of this function returns  $\text{Antecedent}(lit)$ . The other functions and variables are self-explanatory.

`ANALYZE-CONFLICT` progresses from right to left on the conflict graph, starting from the conflicting clause, while constructing the new conflict clause through a series of resolution steps. It begins with the conflicting clause  $cl$ , in which all literals are set to 0. The literal  $lit$  is the literal in  $cl$  assigned last, and  $var$  denotes its associated variable. The antecedent clause of  $var$ , denoted by  $ante$ , contains  $\neg lit$  as the only satisfied literal, and other literals, all of which are currently unsatisfied. The clauses  $cl$  and  $ante$  thus contain  $lit$  and  $\neg lit$ , respectively, and can therefore be resolved with the resolution variable  $var$ . The resolvent clause is again a conflicting clause, which is the basis for the next resolution step.

**Example 2.12.** Consider the partial implication graph and set of clauses in Fig. 2.9, and assume that the implication order in the BCP was  $x_4, x_5, x_6, x_7$ .

The conflict clause  $c_5 := (x_{10} \vee x_2 \vee \neg x_4)$  is computed through a series of binary resolutions. `ANALYZE-CONFLICT` traverses backwards through the implication graph starting from the conflicting clause  $c_4$ , while following the order of the implications in reverse, as can be seen in the table below. The intermediate clauses, in this case the second and third clauses in the resolution sequence, are typically discarded.



**Fig. 2.9.** A partial implication graph and a set of clauses that demonstrate Algorithm 2.2.2. The first UIP is  $x_4$ , and, correspondingly, the asserted literal is  $\neg x_4$

name	$cl$	$lit$	$var$	$ante$
$c_4$	$(\neg x_6 \vee x_7)$	$x_7$	$x_7$	$c_3$
	$(\neg x_5 \vee \neg x_6)$	$\neg x_6$	$x_6$	$c_2$
	$(\neg x_4 \vee x_{10} \vee \neg x_5)$	$\neg x_5$	$x_5$	$c_1$
$c_5$	$(\neg x_4 \vee x_2 \vee x_{10})$			

The clause  $c_5$  is an asserting clause in which the negation of the first UIP ( $x_4$ ) is the only literal from the current decision level.  $\blacksquare$

### 2.2.5 Decision Heuristics

Probably the most important element in SAT solving is the strategy by which the variables and the value given to them are chosen. This strategy is called the **decision heuristic** of the SAT solver. Let us survey some of the best-known decision heuristics, in the order in which they were suggested, which is also the order of their average efficiency as measured by numerous experiments. New strategies are published every year.

#### Jeroslow–Wang

Given a CNF formula  $\mathcal{B}$ , compute for each literal  $l$

$$J(l) = \sum_{\omega \in \mathcal{B}, l \in \omega} 2^{-|\omega|}, \quad (2.10)$$

where  $\omega$  represents a clause and  $|\omega|$  its length. Choose the literal  $l$  for which  $J(l)$  is maximal, and for which neither  $l$  or  $\neg l$  is asserted.

This strategy gives higher priority to literals that appear frequently in short clauses. It can be implemented statically (one computation in the beginning of the run) or dynamically, where in each decision only unsatisfied clauses are considered in the computation. In the context of a SAT solver that learns through addition of conflict clauses, the dynamic approach is more reasonable.

### Dynamic Largest Individual Sum (DLIS)

At each decision level, choose the unassigned literal that satisfies the largest number of currently unsatisfied clauses.

The common way to implement such a heuristic is to keep a pointer from each literal to a list of clauses in which it appears. At each decision level, the solver counts the number of clauses that include this literal and are not yet satisfied, and assigns this number to the literal. Subsequently, the literal with the largest count is chosen. DLIS imposes a large overhead, since the complexity of making a decision is proportional to the number of clauses. Another variation of this strategy, suggested by Copty et al. [52], is to count the number of satisfied clauses resulting from each possible decision *and its implications through BCP*. This variation indeed makes better decisions, but also imposes more overhead.

### Variable State Independent Decaying Sum (VSIDS)

This is a strategy similar to DLIS, with two differences. First, when counting the number of clauses in which every literal appears, we disregard the question of whether that clause is already satisfied or not. This means that the estimation of the quality of every decision is compromised, but the complexity of making a decision is better: it takes a constant time to make a decision assuming we keep the literals in a list sorted by their score. Second, we periodically divide all scores by 2.

The idea is to make the decision heuristic **conflict-driven**, which means that it tries to solve conflicts before attempting to satisfy more original clauses. For this purpose, it needs to give higher scores to variables that are involved in recent conflicts. Recall that every conflict results in a conflict clause. A new conflict clause, like any other clause, adds 1 to the score of each literal that appears in it. The greater the amount of time that has passed since this clause was added, the more often the score of these literals is divided by 2. Thus, variables in new conflict clauses become more influential. The SAT solver CHAFF, which introduced VSIDS, allows one to tune this strategy by controlling the frequency with which the scores are divided and the constant by which they are divided. It turns out that different families of CNF formulas are best solved with different parameters.

### Berkmin

Maintain a score per variable, similar to the score VSIDS maintains for each literal (i.e., increase the counter of a variable if one of its literals appears in a clause, and periodically divide the counters by a constant). Maintain a similar score for each literal, but do not divide it periodically. Push conflict clauses into a stack. When a decision has to be made, search for the topmost clause on this stack that is unresolved. From this clause, choose the unassigned

variable with the highest variable score. Determine the value of this variable by choosing the literal corresponding to this variable with the highest literal score. If the stack is empty, the same strategy is applied, except that the variable is chosen from the set of all unassigned variables rather than from a single clause.

This heuristic was first implemented in a SAT solver called BERKMIN. The idea is to give variables that appear in recent conflicts absolute priority, which seems empirically to be more effective. It also concentrates only on unresolved conflicts, in contrast to VSIDS.

### 2.2.6 The Resolution Graph and the Unsatisfiable Core

Since each conflict clause is derived from a set of other clauses, we can keep track of this process with a **resolution graph**.

**Definition 2.13 (binary resolution graph).** *A binary resolution graph is a directed acyclic graph where each node is labeled with a clause, each root corresponds to an original clause, and each nonroot node has exactly two incoming edges and corresponds to a clause derived by binary resolution from its parents in the graph.*

Typically, SAT solvers do not retain all the intermediate clauses that are created during the resolution process of the conflict clause. They store enough clauses, however, for building a graph that describes the relation between the conflict clauses.


**Definition 2.14 (resolution graph).** *A resolution graph is a directed acyclic graph where each node is labeled with a clause, each root corresponds to an original clause, and each nonroot node has two or more incoming edges and corresponds to a clause derived by resolution from its parents in the graph, possibly through other clauses that are not represented in the graph.*

Resolution graphs are also called **hyperresolution graphs**, to emphasize that they are not necessarily binary.

**Example 2.15.** Consider once again the implication graph in Fig. 2.9. The clauses  $c_1, \dots, c_4$  participate in the resolution of  $c_5$ . The corresponding resolution graph appears in Fig. 2.10. ▀

In the case of an unsatisfiable formula, the resolution graph has a sink node (i.e., a node with incoming edges only), which corresponds to an empty clause.<sup>4</sup>

<sup>4</sup> In practice, SAT solvers terminate before they actually derive the empty clause, as can be seen in Algorithms 2.2.1 and 2.2.2, but it is possible to continue developing the resolution graph after the run is over and derive a full resolution proof ending with the empty clause.



**Fig. 2.10.** A resolution graph corresponding to the implication graph in Fig. 2.9

The resolution graph can be used for various purposes, some of which we mention here. The most common use of this graph is for deriving an *unsatisfiable core* of unsatisfiable formulas.

**Definition 2.16 (unsatisfiable core).** *An unsatisfiable core of a CNF unsatisfiable formula is any unsatisfiable subset of the original set of clauses.*

Unsatisfiable cores which are relatively small subsets of the original set of clauses are useful in various contexts, because they help us to focus on a *cause* of unsatisfiability (there can be multiple unsatisfiable cores not contained in each other, and not even intersecting each other). We leave it to the reader in Problem 2.13 to find an algorithm that computes a core given a resolution graph.

Another common use of a resolution graph is for certifying a SAT solver’s conclusion that a formula is unsatisfiable. Unlike the case of satisfiable instances, for which the satisfying assignment is an easy-to-check piece of evidence, checking an unsatisfiability result is harder. Using the resolution graph, however, an independent checker can replay the resolution steps starting from the original clauses until it derives the empty clause. This verification requires time that is linear in the size of the resolution proof.

### 2.2.7 SAT Solvers: Summary

In this section we have covered the basic elements of modern DPLL solvers, including decision heuristics, learning with conflict clauses, and conflict-driven backtracking. There are various other mechanisms for gaining efficiency that we do not cover in this book, such as efficient implementation of BCP, detection of subsumed clauses, preprocessing and simplification of the formula, deletion of conflict clauses, and **restarts** (i.e., restarting the solver when it seems to be in a hopeless branch of the search tree). The interested reader is referred to the references given in Sect. 2.5.

Let us now reflect on the two approaches to formal reasoning that we described in Sect. 1.1 – deduction and enumeration. Can we say that SAT solvers, as described in this section, follow either one of them? On the one hand, SAT solvers can be thought of as searching a binary tree with  $2^n$  leaves,



where  $n$  is the number of Boolean variables in the input formula. Every leaf is a full assignment, and, hence, traversing all leaves corresponds to enumeration. From this point of view, conflict clauses are generated in order to prune the search space. On the other hand, conflict clauses are *deduced* via the resolution rule from other clauses. If the formula is unsatisfiable then the sequence of applications of this rule, as listed in the SAT solver’s log, is a legitimate deductive proof of unsatisfiability. The search heuristic can therefore be understood as a strategy of applying an inference rule. Thus, the two points of view are equally legitimate.

## 2.3 Binary Decision Diagrams

### 2.3.1 From Binary Decision Trees to ROBDDs

Reduced ordered **binary decision diagrams** (ROBDDs, or **BDDs** for short), are a highly useful graph-based data structure for manipulating Boolean formulas. Unlike CNF, this data representation is **canonical**, which means that if two formulas are equivalent, then their BDD representations are equivalent as well (to achieve this property the two BDDs should be constructed following the same variable order, as we will soon explain). Canonicity is *not* a property of CNF, DNF, or NNF (see Sect. 1.3). Consider, for example, the two CNF formulas

$$\mathcal{B}_1 := (x_1 \wedge (x_2 \vee x_3)), \quad \mathcal{B}_2 := (x_1 \wedge (x_1 \vee x_2) \wedge (x_2 \vee x_3)). \quad (2.11)$$


Although the two formulas are in the same normal form and logically equivalent, they are syntactically different. The BDD representations of  $\mathcal{B}_1$  and  $\mathcal{B}_2$ , on the other hand, are the same.

One implication of canonicity is that all tautologies have the same BDD (a single node with a label “1”) and all contradictions also have the same BDD (a single node with a label “0”). Thus, although two CNF formulas of completely different size can both be unsatisfiable, their BDD representations are identical: a single node with the label “0”. As a consequence, checking for satisfiability, validity, or contradiction can be done in constant time for a given BDD. There is no free lunch, however: building the BDD for a given formula can take exponential space and time, even if in the end it results in a single node.

We start with a simple **binary decision tree** to represent a Boolean formula. Consider the formula

$$\mathcal{B} := ((x_1 \wedge x_2) \vee (\neg x_1 \wedge x_3)). \quad (2.12)$$

The binary decision tree in Fig. 2.11 represents this formula with the variable ordering  $x_1, x_2, x_3$ . Notice how this order is maintained in each path along the




**Fig. 2.11.** A binary decision tree for (2.12). The drawing follows the convention by which dashed edges represent an assignment of 0 to the variable labeling the source node


tree, and that each of these variables appears exactly once in each path from the root to one of the leaves.

Such a binary decision tree is not any better, in terms of space consumption, than an explicit truth table, as it has  $2^n$  leaves. Every path in this tree, from root to leaf, corresponds to an assignment. Every path that leads to a leaf “1” corresponds to a *satisfying* assignment. For example, the path  $x_1 = 1, x_2 = 1, x_3 = 0$  corresponds to a satisfying assignment of our formula  $\mathcal{B}$  because it ends in a leaf with the label “1”. Altogether, four assignments satisfy this formula. The question is whether we can do better than a binary decision tree in terms of space consumption, as there is obvious redundancy in this tree. We now demonstrate the three **reduction rules** that can be applied to such trees. Together they define what a *reduced ordered BDD* is.

- **Reduction #1.** *Merge the leaf nodes into two nodes “1” and “0”.* The result of this reduction appears in Fig. 2.12.
- **Reduction #2.** *Merge isomorphic subtrees.* Isomorphic subtrees are subtrees that have roots that represent the same variable (if these are leaves, then they represent the same Boolean value), and have left and right children that are isomorphic as well. After applying this rule to our graph, we are left with the diagram in Fig. 2.13. Note how the subtrees rooted at the left two nodes labeled with  $x_3$  are isomorphic and are therefore merged in this reduction.
- **Reduction #3.** *Removing redundant nodes.* In the diagram in Fig. 2.13, it is clear that the left  $x_2$  node is redundant, because its value does not affect the values of paths that go through it. The same can be said about the middle and right nodes corresponding to  $x_3$ . In each such case, we can simply remove the node, while redirecting its incoming edge to the node



**Fig. 2.12.** After applying reduction #1, merging the leaf nodes into two nodes



**Fig. 2.13.** After applying reduction #2, merging isomorphic subtrees

to which both of its edges point. This reduction results in the diagram in Fig. 2.14.

The second and third reductions are repeated as long as they can be applied. At the end of this process, the BDD is said to be *reduced*.

Several important properties of binary trees are maintained during the reduction process:


1. Each terminal node  $v$  is associated with a Boolean value  $val(v)$ . Each nonterminal node  $v$  is associated with a variable, denoted by  $var(v) \in Var(\mathcal{B})$ .
2. Every nonterminal node  $v$  has exactly two children, denoted by  $low(v)$  and  $high(v)$ , corresponding to a FALSE or TRUE assignment to  $var(v)$ .
3. Every path from the root to a leaf node contains not more than one occurrence of each variable. Further, the order of variables in each such path is consistent with the order in the original binary tree.

$val(v)$

$var(v)$

$low(v)$

$high(v)$



**Fig. 2.14.** After applying reduction #3, removing redundant nodes

4. A path to the “1” node through all variables corresponds to an assignment that satisfies the formula.

Unlike a binary tree, a BDD can have paths to the leaf nodes through only *some* of the variables. Such paths to the “1” node satisfy the formula regardless of the values given to the other variables, which are appropriately known by the name **don’t cares**. A reduced BDD has the property that it does not contain any redundant nodes or isomorphic subtrees, and, as indicated earlier, it is canonical.


### 2.3.2 Building BDDs from Formulas

The process of turning a binary tree into a BDD helps us to explain the reduction rules, but is not very useful by itself, as we do not want to build the binary decision tree in the first place, owing to its exponential size. Instead, we create the ROBDDs directly: given a formula, we build its BDD recursively from the BDDs of its subexpressions. For this purpose, Bryant defined the procedure `APPLY`, which, given two BDDs  $\mathcal{B}$  and  $\mathcal{B}'$ , builds a BDD for  $\mathcal{B} \star \mathcal{B}'$ , where  $\star$  stands for any one of the 16 binary Boolean operators (such as “ $\wedge$ ”, “ $\vee$ ”, and “ $\implies$ ”). The complexity of `APPLY` is bounded by  $|\mathcal{B}| \cdot |\mathcal{B}'|$ , where  $|\mathcal{B}|$  and  $|\mathcal{B}'|$  denote the respective sizes of  $\mathcal{B}$  and  $\mathcal{B}'$ .

In order to describe `APPLY`, we first need to define the **restrict** operation. This operation is simply an assignment of a value to one of the variables in the BDD. We denote the restriction of  $\mathcal{B}$  to  $x = 0$  by  $\mathcal{B}|_{x=0}$  or, in other words, the BDD corresponding to the function  $\mathcal{B}$  after assigning 0 to  $x$ . Given the BDD for  $\mathcal{B}$ , it is straightforward to compute its restriction to  $x = 0$ . For every node  $v$  such that  $\text{var}(v) = x$ , we remove  $v$  and redirect the incoming edges of  $v$  to  $\text{low}(v)$ . Similarly, if the restriction is  $x = 1$ , we redirect all the incoming edges to  $\text{high}(v)$ .

$\mathcal{B} \star \mathcal{B}'$

$\mathcal{B}|_{x=0}$



**Fig. 2.15.** Restricting  $\mathcal{B}$  to  $x_2 = 0$ . This operation is denoted by  $\mathcal{B}|_{x_2=0}$

Let  $\mathcal{B}$  denote the function represented by the BDD in Fig. 2.14. The diagram in Fig. 2.15 corresponds to  $\mathcal{B}|_{x_2=0}$ , which is the function  $\neg x_1 \wedge x_3$ . Let  $v$  and  $v'$  denote the root variables of  $\mathcal{B}$  and  $\mathcal{B}'$ , respectively, and let  $\text{var}(v) = x$  and  $\text{var}(v') = x'$ . APPLY operates recursively on the BDD structure, following one of these four cases:

1. If  $v$  and  $v'$  are both terminal nodes, then  $\mathcal{B} \star \mathcal{B}'$  is a terminal node with the value  $\text{val}(v) \star \text{val}(v')$ .
2. If  $x = x'$ , that is, the roots of both  $\mathcal{B}$  and  $\mathcal{B}'$  correspond to the same variable, then we apply what is known as **Shannon expansion**:

$$\mathcal{B} \star \mathcal{B}' := (\neg x \wedge (\mathcal{B}|_{x=0} \star \mathcal{B}'|_{x=0})) \vee (x \wedge (\mathcal{B}|_{x=1} \star \mathcal{B}'|_{x=1})). \quad (2.13)$$

Thus, the resulting BDD has a new node  $v''$  such that  $\text{var}(v'') = x$ ,  $\text{low}(v'')$  points to a BDD representing  $\mathcal{B}|_{x=0} \star \mathcal{B}'|_{x=0}$ , and  $\text{high}(v'')$  points to a BDD representing  $\mathcal{B}|_{x=1} \star \mathcal{B}'|_{x=1}$ . Note that both of these restricted BDDs refer to a smaller set of variables than do  $\mathcal{B}$  and  $\mathcal{B}'$ . Therefore, if  $\mathcal{B}$  and  $\mathcal{B}'$  refer to the same set of variables, then this process eventually reaches the leaves, which are handled by the first case.

3. If  $x \neq x'$  and  $x$  precedes  $x'$  in the given variable order, we again apply Shannon expansion, except that this time we use the fact that the value of  $x$  does not affect the value of  $\mathcal{B}'$ , that is,  $\mathcal{B}'|_{x=0} = \mathcal{B}'|_{x=1} = \mathcal{B}'$ . Thus, the formula above simplifies to


$$\mathcal{B} \star \mathcal{B}' := (\neg x \wedge (\mathcal{B}|_{x=0} \star \mathcal{B}')) \vee (x \wedge (\mathcal{B}|_{x=1} \star \mathcal{B}')). \quad (2.14)$$

Once again, the resulting BDD has a new node  $v''$  such that  $\text{var}(v'') = x$ ,  $\text{low}(v'')$  points to a BDD representing  $\mathcal{B}|_{x=0} \star \mathcal{B}'$ , and  $\text{high}(v'')$  points to a BDD representing  $\mathcal{B}|_{x=1} \star \mathcal{B}'$ . Thus, the only difference is that we reuse  $\mathcal{B}'$  in the recursive call as is, instead of its restriction to  $x = 0$  or  $x = 1$ .


4. The case in which  $x \neq x'$  and  $x$  follows  $x'$  in the given variable order is dual to the previous case.

We now demonstrate APPLY with an example.

**Example 2.17.** Assume that we are given the BDDs for  $\mathcal{B} := (x_1 \iff x_2)$  and for  $\mathcal{B}' := \neg x_2$ , and that we want to compute the BDD for  $\mathcal{B} \vee \mathcal{B}'$ . Both the source BDDs and the target BDD follow the same order:  $x_1, x_2$ . Figure 2.16 presents the BDDs for  $\mathcal{B}$  and  $\mathcal{B}'$ .




**Fig. 2.16.** The two BDDs corresponding to  $\mathcal{B} := (x_1 \iff x_2)$  (left) and  $\mathcal{B}' := \neg x_2$  (right)




**Fig. 2.17.** Since  $x_1$  appears before  $x_2$  in the variable order, we apply case 3

Since the root nodes of the two BDDs are different, we apply case 3. This results in the diagram in Fig. 2.17. In order to compute the BDD for  $\mathcal{B}|_{x_1=0} \vee \mathcal{B}'$ , we first compute  $\mathcal{B}|_{x_1=0}$ . This results in the diagram on the left of Fig. 2.18. To compute  $\mathcal{B}|_{x_1=0} \vee \mathcal{B}'$ , we apply case 2, as the root nodes refer to the same variable,  $x_2$ . This results in the BDD on the right of the figure. Repeating the same process for  $high(x_1)$ , results in the leaf BDD “1”, and thus our final BDD is as shown in Fig. 2.19. This BDD represents the function  $x_1 \vee (\neg x_1 \wedge \neg x_2)$ , which is indeed the result of  $\mathcal{B} \vee \mathcal{B}'$ . ▀

The size of the BDD depends strongly on the variable order. That is, constructing the BDD for a given function using different variable orders results in radically different BDDs. There are functions for which one BDD order results in a BDD with a polynomial number of nodes, whereas with a different order the number of nodes is exponential. Bryant gives the function



**Fig. 2.18.** Applying case 2, since the root nodes refer to the same variable. The left and right leaf nodes of the resulting BDD are computed following case 1, since the source nodes are leaves



**Fig. 2.19.** The final BDD for  $\mathcal{B} \vee \mathcal{B}'$

$(x_1 \iff x'_1) \wedge \dots \wedge (x_n \iff x'_n)$  as an example of this phenomenon: using the variable order  $x_1, x'_1, x_2, x'_2, \dots, x_n, x'_n$ , the size of the BDD is  $3n + 2$  while with the order  $x_1, x_2, \dots, x_n, x'_1, x'_2, \dots, x'_n$ , the BDD has  $3 \cdot 2^n - 1$  nodes. Furthermore, there are functions for which there is no variable order that results in a polynomial number of nodes. Multiplication of bit vectors (arrays of Boolean variables; see Chap. 6) is one such well-known example. Finding a good variable order is a subject that has been researched extensively and has yielded many PhD theses. It is an NP-complete problem to decide whether a given variable order is optimal [36]. Recall that once the BDD has been built, checking satisfiability and validity is a constant-time operation. Thus, if we could always easily find an order in which building the BDD takes polynomial time, this would make satisfiability and validity checking a polynomial-time operation.

There is a very large body of work on BDDs and their extensions – variable-ordering strategies is only one part of this work. Extending BDDs to handle variables of types other than Boolean is an interesting subject, which we briefly discuss as part of Problem 2.15. Another interesting topic is alternatives to APPLY. As part of Problem 2.14, we describe one such alternative based on a recursive application of the *ite* (if-then-else) function.

## 2.4 Problems

### 2.4.1 Warm-up Exercises

**Problem 2.1 (modeling: simple).** Consider three persons A, B, and C who need to be seated in a row. But:

- A does not want to sit next to C.
- A does not want to sit in the left chair.
- B does not want to sit to the right of C.

Write a propositional formula that is satisfiable if and only if there is a seat assignment for the three persons that satisfies all constraints. Is the formula satisfiable? If so, give an assignment.

**Problem 2.2 (modeling: program equivalence).** Show that the two if-then-else expressions below are equivalent:

$$!(a \parallel b) ? h : !(a == b) ? f : g \quad !(a \parallel !b) ? g : (!a \&\& !b) ? h : f$$

You can assume that the variables have only one bit.

**Problem 2.3 (SAT solving).** Consider the following set of clauses:

$$\begin{aligned} &(x_5 \vee \neg x_1 \vee x_3), (\neg x_1 \vee x_2), \\ &(\neg x_2 \vee x_4), (\neg x_3 \vee \neg x_4), \\ &(\neg x_5 \vee x_1), (\neg x_5 \vee \neg x_6), \\ &(x_6 \vee x_1). \end{aligned} \tag{2.15}$$

Apply the Berkmin decision heuristic, including the application of ANALYZE-CONFLICT with conflict-driven backtracking. In the case of a tie (during the application of VSIDS), make a decision that leads to a conflict. Show the implication graph at each decision level.

**Problem 2.4 (BDDs).** Construct the BDD for  $\neg(x_1 \vee (x_2 \wedge \neg x_3))$  with the variable order  $x_1, x_2, x_3$ ,

- starting from a decision tree, and
- bottom-up (starting from the BDDs of the atoms  $x_1, x_2, x_3$ ).

### 2.4.2 Modeling

**Problem 2.5 (unwinding a finite automaton).** A *nondeterministic finite automaton* is a 5-tuple  $\langle Q, \Sigma, \delta, I, F \rangle$ , where

- $Q$  is a finite set of states,
- $\Sigma$  is the alphabet (a finite set of letters),
- $\delta : Q \times \Sigma \longrightarrow 2^Q$  is the transition function ( $2^Q$  is the power set of  $Q$ ),
- $I \subseteq Q$  is the set of initial states, and




- $F \subseteq Q$  is the set of accepting states.

The transition function determines to which states we can move given the current state and input. The automaton is said to *accept* a finite input string  $s_1, \dots, s_n$  with  $s_i \in \Sigma$  if and only if there is a sequence of states  $q_0, \dots, q_n$  with  $q_i \in Q$  such that

- $q_0 \in I$ ,
- $\forall i \in \{1, \dots, n\}. q_i \in \delta(q_{i-1}, s_i)$ , and
- $q_n \in F$ .

For example, the automaton in Fig. 2.20 is defined by  $Q = \{s_1, s_2\}$ ,  $\Sigma = \{a, b\}$ ,  $\delta(s_1, a) = \{s_1\}$ ,  $\delta(s_1, b) = \{s_1, s_2\}$ ,  $I = \{s_1\}$ ,  $F = \{s_2\}$ , and accepts strings that end with  $b$ . Given a nondeterministic finite automaton  $\langle Q, \Sigma, \delta, I, F \rangle$  and a fixed input string  $s_1, \dots, s_n$ ,  $s_i \in \Sigma$ , construct a propositional formula that is satisfiable if and only if the automaton accepts the string.



**Fig. 2.20.** A nondeterministic finite automaton accepting all strings ending with the letter  $b$

**Problem 2.6 (assigning teachers to subjects).** A problem of covering  $m$  subjects with  $k$  teachers may be defined as follows. Let  $T : \{T_1, \dots, T_n\}$  be a set of teachers. Let  $S : \{S_1, \dots, S_m\}$  be a set of subjects. Each teacher  $t \in T$  can teach some subset  $S(t)$  of the subjects  $S$  (i.e.,  $S(t) \subseteq S$ ). Given a natural number  $k \leq n$ , is there a subset of size  $k$  of the teachers that together covers all  $m$  subjects, i.e., a subset  $C \subseteq T$  such that  $|C| = k$  and  $(\bigcup_{t \in C} S(t)) = S$ ?

**Problem 2.7 (Hamiltonian cycle).** Show a formulation in propositional logic of the following problem: given a directed graph, does it contain a Hamiltonian cycle (a closed path that visits each node, other than the first, exactly once)?

### 2.4.3 Complexity

**Problem 2.8 (space complexity of DPLL with learning).** What is the worst-case space complexity of a DPLL SAT solver as described in Sect. 2.2, in the following cases

- Without learning,
- With learning, i.e., by recording conflict clauses,
- With learning in which the length of the recorded conflict clauses is bounded by a natural number  $k$ .

**Problem 2.9 (polynomial-time (restricted) SAT).** Consider the following two restriction of CNF:

- A CNF in which there is not more than one positive literal in each clause.
  - A CNF formula in which no clause has more than two literals.
1. Show a polynomial-time algorithm that solves each of the problems above.
  2. Show that every CNF can be converted to another CNF which is a conjunction of the two types of formula above. In other words, in the resulting formula all the clauses are either unary, binary, or have not more than one positive literal. How many additional variables are necessary for the conversion?

#### 2.4.4 DPLL SAT Solving

**Problem 2.10 (backtracking level).** We saw that SAT solvers working with conflict-driven backtracking backtrack to the second highest decision level  $dl$  in the asserting conflict clause. This wastes all of the work done from decision level  $dl + 1$  to the current one, say  $dl'$  (although, as we mentioned, this has other advantages that outweigh this drawback). Suppose we try to avoid this waste by performing conflict-driven backtracking as usual, but then repeat the assignments from levels  $dl + 1$  to  $dl' - 1$  (i.e., override the standard decision heuristic for these decisions). Can it be guaranteed that this reassignment will progress without a conflict?

**Problem 2.11 (is the first UIP well defined?).** Prove that in a conflict graph, the notion of a first UIP is well defined, i.e., there is always a single UIP closest to the conflict node. Hint: you may use the notion of *dominators* from graph theory.

#### 2.4.5 Related Problems

**Problem 2.12 (incremental satisfiability).** Given two CNF formulas  $C_1$  and  $C_2$ , under what conditions can a conflict clause learned while solving  $C_1$  be reused when solving  $C_2$ ? In other words, if  $c$  is a conflict clause learned while solving  $C_1$ , under what conditions is  $C_2$  satisfiable if and only if  $C_2 \wedge c$  is satisfiable? How can the condition that you suggest be implemented inside a SAT solver? *Hint:* think of CNF formulas as sets of clauses.

**Problem 2.13 (unsatisfiable cores).**

- (a) Suggest an algorithm that, given a resolution graph (see Definition 2.14), finds an unsatisfiable core of the original formula that is small as possible (by this we do not mean that it has to be minimal).
- (b) Given an unsatisfiable core, suggest a method that attempts to minimize it further.

### 2.4.6 Binary Decision Diagrams

**Problem 2.14 (implementing APPLY with *ite*).** (Based on [29]) Efficient implementations of BDD packages do not use APPLY; rather they use a recursive procedure based on the *ite* (if-then-else) operator. All binary Boolean operators can be expressed as such expressions. For example,

$$\begin{aligned} f \vee g &= ite(f, 1, g), & f \wedge g &= ite(f, g, 0), \\ f \oplus g &= ite(f, \neg g, g), & \neg f &= ite(f, 0, 1). \end{aligned} \quad (2.16)$$

How can a BDD for the *ite* operator be constructed? Assume that  $x$  labels the root nodes of two BDDs  $f$  and  $g$ , and that we need to compute  $ite(c, f, g)$ . Observe the following equivalence:

$$ite(c, f, g) = ite(x, ite(c|_{x=1}, f|_{x=1}, g|_{x=1}), ite(c|_{x=0}, f|_{x=0}, g|_{x=0})). \quad (2.17)$$

Hence, we can construct the BDD for  $ite(c, f, g)$  on the basis of a recursive construction. The root node of the result is  $x$ ,  $low(x) = ite(c|_{x=0}, f|_{x=0}, g|_{x=0})$ , and  $high(x) = ite(c|_{x=1}, f|_{x=1}, g|_{x=1})$ . The terminal cases are

$$\begin{aligned} ite(1, f, g) &= ite(0, g, f) = ite(f, 1, 0) = ite(g, f, f) = f, \\ ite(f, 0, 1) &= \neg f. \end{aligned}$$

1. Let  $f := (x \wedge y)$ ,  $g := \neg x$ . Show an *ite*-based construction of  $f \vee g$ .
2. Present pseudocode for constructing a BDD for the *ite* operator. Describe the data structure that you assume. Explain how your algorithm can be used to replace APPLY.


**Problem 2.15 (binary decision diagrams for non-Boolean functions).** (Based on [47].) Let  $f$  be a function mapping a vector of  $m$  Boolean variables to an integer, i.e.,  $f : B^m \mapsto \mathbb{Z}$ , where  $B = \{0, 1\}$ .

Let  $\{I_1, \dots, I_N\}$ ,  $N \leq 2^m$ , be the set of possible values of  $f$ . The function  $f$  partitions the space  $B^m$  of Boolean vectors into  $N$  sets  $\{S_1, \dots, S_N\}$ , such that for  $i \in \{1 \dots N\}$ ,  $S_i = \{\bar{x} \mid f(\bar{x}) = I_i\}$  (where  $\bar{x}$  denotes a vector). Let  $f_i$  be a characteristic function of  $S_i$  (i.e., a function mapping a vector  $\bar{x}$  to 1 if  $f(\bar{x}) \in S_i$  and to 0 otherwise). Every function  $f(\bar{x})$  can be rewritten as  $\sum_{i=1}^N f_i(\bar{x}) \cdot I_i$ , a form that can be represented as a BDD with  $\{I_1, \dots, I_N\}$  as its terminal nodes. Figure 2.21 shows such a *multiterminal binary decision diagram* (MTBDD) for the function  $2x_1 + 2x_2$ .

Show an algorithm for computing  $f \odot g$ , where  $f$  and  $g$  are multiterminal BDDs, and  $\odot$  is some arithmetic binary operation. Compute with your algorithm the MTBDD of  $f \odot g$ , where

$$\begin{aligned} f &:= \text{if } x_1 \text{ then } 2x_2 + 1 \text{ else } -x_2, \\ g &:= \text{if } x_2 \text{ then } 4x_1 \text{ else } x_3 + 1, \end{aligned}$$

following the variable order  $x_1, x_2, x_3$ .



**Fig. 2.21.** A multiterminal BDD for the function  $f(x, y) = 2x_1 + 2x_2$

## 2.5 Bibliographic Notes

### SAT

The Davis–Putnam–Loveland–Logemann framework was a two-stage invention. In 1960, Davis and Putnam considered CNF formulas and offered a procedure to solve it based on an iterative application of three rules [57]: the pure literal rule, the unit clause rule (what we now call BCP), and what they called “the elimination rule”, which is a rule for eliminating a variable by invoking resolution (e.g., to eliminate  $x$  from a given CNF, apply resolution to each pair of clauses of the form  $(x \vee A) \wedge (\neg x \vee B)$ , erase the resolving clauses, and maintain the resolvent). Their motivation was to optimize a previously known incomplete technique for deciding first-order formulas. Note that at the time, “optimizing” also meant a procedure that was easier to conduct by hand. In 1962, Loveland and Logemann, two programmers hired by Davis and Putnam to implement their idea, concluded that it was more efficient to split and backtrack rather than to apply resolution, and together with Davis published what we know today as the basic DPLL framework [56]. Numerous SAT solvers were developed through the years on the basis of this framework. The alternative approach of stochastic solvers, which were not discussed in length in this chapter, was led for many years by the GSAT and WALKSAT solvers [176].

The definition of the constraints satisfaction problem (CSP) [132] by Montanari (and even before that by Waltz in 1975), a problem which generalizes SAT to arbitrary finite discrete domains and arbitrary constraints, and the development of efficient CSP solvers, led to cross-fertilization between the two fields: nonchronological backtracking, for example, was first used with the CSP, and then adopted by Marques-Silva and Sakallah for their GRASP SAT solver [182], which was the fastest from 1996 to 2000. The addition of conflict clauses in GRASP was also influenced (although in significantly changed form) by earlier techniques called *no-good recording* that were applied to CSP solvers. Bayardo and Schrag [15] also published a method for adapting conflict-driven learning to SAT. The introduction of CHAFF in 2001 [133]

by Moskewicz, Madigan, Zhao, Zhang and Malik marked a breakthrough in performance that led to renewed interest in the field. These authors introduced the idea of conflict-driven nonchronological backtracking coupled with VSIDS, the first conflict-driven decision heuristic. They also introduced a new mechanism for performing fast BCP, a subject not covered in this chapter, empirically identified the first UIP scheme as the most efficient out of various alternative schemes, and introduced many other means for efficiency. The solver SIEGE introduced Variable-Move-To-Front (VMTF), a decision heuristic that moves a constant number of variables from the conflict clause to the top of the list, which performs very well in practice [171]. An indication of how rapid the progress in this field has been was given in the 2006 SAT competition: the best solver in the 2005 competition took ninth place, with a large gap in the run time compared with the 2006 winner, MINISAT-2 [73]. New SAT solvers are introduced every year; readers interested in the latest tools should check the results of the annual SAT competitions. In 2007 the solver RSAT [151] won the “industrial benchmarks” category. RSAT was greatly influenced by MINISAT, but includes various improvements such as ordering of the implications in the BCP stack, an improved policy for restarting the solver, and repeating assignments that are erased while backtracking.

The realization that different classes of problems (e.g., random instances, industrial instances from various problem domains, crafted problems) are best solved with different solvers (or different run time parameters of the same solvers), led to a strategy of invoking an **algorithm portfolio**. This means that one out of  $n$  predefined solvers is chosen automatically for a given problem instance, based on a prediction of which solver is likely to perform best. First, a large “training set” is used for building **empirical hardness models** [143] based on various attributes of the instances in this set. Then, given a problem instance, the run time of each of the  $n$  solvers is predicted, and accordingly the solver is chosen for the task. SATzilla [205] is a successful algorithm portfolio based on these ideas that won several categories in the 2007 competition.

Zhang and Malik described a procedure for efficient extraction of unsatisfiable cores and unsatisfiability proofs from a SAT solver [210, 211]. There are many algorithms for minimizing such cores – see, for example, [81, 98, 118, 144]. The description of the main SAT procedure in this chapter was inspired mainly by [210, 211]. BERKMIN, a SAT solver developed by Goldberg and Novikov, introduced what we have named “the Berkmin decision heuristic” [88]. The connection between the process of deriving conflict clauses and resolution was discussed in, for example, [16, 80, 116, 207, 210].

Incremental satisfiability in its modern version, i.e., the problem of which conflict clauses can be reused when solving a related problem (see Problem 2.12) was introduced by Strichman in [180, 181] and independently by Whitemore, Kim, and Sakallah in [197]. Earlier versions of this problem were more restricted, for example the work of Hooker [96] and of Kim, Whitemore, Marques-Silva, and Sakallah [105].

There is a large body of theoretical work on SAT as well. Probably the best-known is related to complexity theory: SAT played a major role in the theoretical breakthrough achieved by Cook in 1971 [50], who showed that every NP problem can be reduced to SAT. Since SAT is in NP, this made it the first problem to be identified as belonging to the NP-complete complexity class. The general scheme for these reductions (through a translation to a Turing machine) is rarely used and is not efficient. Direct translations of almost all of the well-known NP problems have been suggested through the years, and, indeed, it is always an interesting question whether it is more efficient to solve problems directly or to reduce them to SAT (or to any other NP-complete problem, for that matter). The incredible progress in the efficiency of these solvers in the last decade has made it very appealing to take the translation option. By translating problems to CNF we may lose high-level information about the problem, but we can also gain low-level information that is harder to detect in the original representation of the problem.

An interesting angle of SAT is that it attracts research by physicists!<sup>5</sup> Among other questions, they attempt to solve the **phase transition** problem [45, 128]: why and when does a randomly generated SAT problem (according to some well-defined distribution) become hard to solve? There is a well-known result showing empirically that randomly generated SAT instances are hardest when the ratio between the numbers of clauses and variables is around 4.2. A larger ratio makes the formula more likely to be unsatisfiable, and the more constraints there are, the easier it is to detect the unsatisfiability. A lower ratio has the opposite effect: it makes the formula more likely to be satisfiable and easier to solve. Another interesting result is that as the formula grows, the phase transition sharpens, asymptotically reaching a sharp phase transition, i.e., a threshold ratio such that all formulas above it are unsatisfiable, whereas all formulas beneath it are satisfiable. There have been several articles about these topics in *Science* [106, 127], *Nature* [131] and even *The New York Times* [102].

## Binary Decision Diagrams

Binary decision diagrams were introduced by Lee in 1959 [115], and explored further by Akers [3]. The full potential for efficient algorithms based on the data structure was investigated by Bryant [35]: his key extensions were to use a fixed variable ordering (for canonical representation) and shared subgraphs (for compression). Together they form what we now refer to as reduced-ordered BDDs. Generally ROBDDs are efficient data structures accompanied by efficient manipulation algorithms for the representation of sets and relations. ROBDDs later became a vital component of symbolic *model checking*, a technique that led to the first large-scale use of formal verification techniques in industry (mainly in the field of electronic design automation). Numerous extensions of ROBDDs exist in the literature, some of which extend

---

<sup>5</sup> The origin of this interest is in statistical mechanics.

the logic that the data structure can represent beyond propositional logic, and some adapt it to a specific need. Multiterminal BDDs (also discussed in Problem 2.15), for example, were introduced in [47] to obtain efficient *spectral transforms*, and multiplicative binary moment diagrams (\*BMDs) [37] were introduced for efficient representation of linear functions. There is also a large body of work on variable ordering in BDDs and dynamic variable reordering (ordering of the variables during the construction of the BDD, rather than according to a predefined list).

It is clear that BDDs can be used everywhere SAT is used (in its basic functionality). SAT is typically more efficient, as it does not require exponential space even in the worst case.<sup>6</sup> The other direction is not as simple, because BDDs, unlike CNF, are canonical. Furthermore, finding all solutions to the Boolean formula represented by a BDD is linear in the number of solutions (all paths leading to the “1” node), while worst-case exponential time is needed for each solution of the CNF. There are various extensions to SAT (algorithms for what is known as **all-SAT**, the problem of finding all solutions to a propositional formula) that attempt to solve this problem in practice using a SAT solver.

## 2.6 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$x_i@d$	(SAT) $x_i$ is assigned TRUE at decision level $d$	30
$val(v)$	(BDD) the 0 or 1 value of a BDD leaf node	45
$var(v)$	(BDD) the variable associated with an internal BDD node	45
$low(v)$	(BDD) the node pointed to by node $v$ when $v$ is assigned 0	45
$high(v)$	(BDD) the node pointed to by node $v$ when $v$ is assigned 1	45
$\mathcal{B} \star \mathcal{B}'$	(BDD) $\star$ is any of the 16 binary Boolean operators	46
$\mathcal{B} _{x=0}$	(BDD) simplification of $\mathcal{B}$ after assigning $x = 0$ (also called “restriction”)	46

<sup>6</sup> This characteristic of SAT can be achieved by restricting the number of added conflict clauses. In practice, even without this restriction, memory is rarely the bottleneck.

---

## Equality Logic and Uninterpreted Functions

### 3.1 Introduction

This chapter introduces the **theory of equality**, also known by the name **equality logic**. Equality logic can be thought of as propositional logic where the atoms are equalities between variables over some infinite type or between variables and constants. As an example, the formula  $(y = z \vee (\neg(x = z) \wedge x = 2))$  is a well-formed equality logic formula, where  $x, y, z \in \mathbb{R}$  ( $\mathbb{R}$  denotes the reals). An example of a satisfying assignment is  $\{x \mapsto 2, y \mapsto 2, z \mapsto 0\}$ .

**Definition 3.1 (equality logic).** *An equality logic formula is defined by the following grammar:*

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid \neg \text{formula} \mid (\text{formula}) \mid \text{atom} \\ \text{atom} &: \text{term} = \text{term} \\ \text{term} &: \text{identifier} \mid \text{constant} \end{aligned}$$

where the identifiers are variables defined over a single infinite domain such as the Reals or Integers.<sup>1</sup> Constants are elements from the same domain as the identifiers.

#### 3.1.1 Complexity and Expressiveness

The satisfiability problem for equality logic is NP-complete. We leave the proof of this claim as an exercise (Problem 4.7 in Chap. 4). The fact that both equality logic and propositional logic are NP-complete implies that they can model the same decision problems (with not more than a polynomial difference in the number of variables). Why should we study both, then?

For two main reasons: convenience of modeling, and efficiency. It is more natural and convenient to use equality logic for modeling certain problems

---

<sup>1</sup> The restriction to a single domain (also called a single type or a single *sort*) is not essential. It is introduced for the sake of simplicity of the presentation.



than to use propositional logic, and vice versa. As for efficiency, the high-level structure in the input equality logic formula can potentially be used to make the decision procedure work faster. This information may be lost if the problem is modeled directly in propositional logic.

### 3.1.2 Boolean Variables

Frequently, equality logic formulas are mixed with Boolean variables. Nevertheless, we shall not integrate them into the definition of the theory, in order to keep the description of the algorithms simple. Boolean variables can easily be eliminated from the input formula by replacing each such variable with an equality between two new variables. But this is not a very efficient solution. As we progress in this chapter, it will be clear that it is easy to handle Boolean variables directly, with only small modifications to the various decision procedures. The same observation applies to many of the other theories that we consider in this book.

### 3.1.3 Removing the Constants: A Simplification

$\varphi^E$  **Theorem 3.2.** *Given an equality logic formula  $\varphi^E$ , there is an algorithm that generates an equisatisfiable formula (see Definition 1.9)  $\varphi^{E'}$  without constants, in polynomial time.*

#### Algorithm 3.1.1: REMOVE-CONSTANTS

**Input:** An equality logic formula  $\varphi^E$  with constants  $c_1, \dots, c_n$

**Output:** An equality logic formula  $\varphi^{E'}$  such that  $\varphi^{E'}$  and  $\varphi^E$  are equisatisfiable and  $\varphi^{E'}$  has no constants

1.  $\varphi^{E'} := \varphi^E$ .
2. In  $\varphi^{E'}$ , replace each constant  $c_i$ ,  $1 \leq i \leq n$ , with a new variable  $C_{c_i}$ .
3. For each pair of constants  $c_i, c_j$  such that  $1 \leq i < j \leq n$ , add the constraint  $C_{c_i} \neq C_{c_j}$  to  $\varphi^{E'}$ .

Algorithm 3.1.1 eliminates the constants from a given formula by replacing them with new variables. Problem 3.2, and, later, Problem 4.4, focus on this procedure. Unless otherwise stated, we assume from here on that the input equality formulas do not have constants.

## 3.2 Uninterpreted Functions

Equality logic is far more useful if combined with **uninterpreted functions**. Uninterpreted functions are used for abstracting, or generalizing, theorems.

Unlike other function symbols, they should not be interpreted as part of a model of a formula. In the following formula, for example,  $F$  and  $G$  are uninterpreted, whereas the binary function symbol “+” is interpreted as the usual addition function:

$$F(x) = F(G(y)) \vee x + 1 = y . \quad (3.1)$$

**Definition 3.3 (equality logic with uninterpreted functions (EUF)).** *An equality logic formula with uninterpreted functions and uninterpreted predicates<sup>2</sup> is defined by the following grammar:*

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid \neg \text{formula} \mid (\text{formula}) \mid \text{atom} \\ \text{atom} &: \text{term} = \text{term} \mid \text{predicate-symbol} (\text{list of terms}) \\ \text{term} &: \text{identifier} \mid \text{function-symbol} (\text{list of terms}) \end{aligned}$$

We generally use capital letters to denote uninterpreted functions, and use the superscript UF to denote EUF formulas.

### Aside: The Logic Perspective

To explain the meaning of uninterpreted functions from the perspective of logic, we have to go back to the notion of a **theory**, which was explained in Sect. 1.4. Recall the set of axioms (1.35), and that in this chapter we refer to the quantifier-free fragment.

Only a single additional axiom (an axiom scheme, actually) is necessary in order to extend equality logic to EUF. For each  $n$ -ary function symbol,  $n > 0$ ,

$$\forall t_1, \dots, t_n, t'_1, \dots, t'_n. \quad \bigwedge_i t_i = t'_i \implies F(t_1, \dots, t_n) = F(t'_1, \dots, t'_n) \quad (\text{CONGRUENCE}), \quad (3.2)$$

where  $t_1, \dots, t_n, t'_1, \dots, t'_n$  should be instantiated with terms that appear as arguments of uninterpreted functions in the formula. A similar axiom can be defined for uninterpreted predicates.

Thus, whereas in theories where the function symbols are interpreted there are axioms to define their semantics – what we want them to mean – in a theory over uninterpreted functions, the only restriction we have over a satisfying interpretation is that imposed by functional consistency, namely the restriction imposed by the (CONGRUENCE) rule.

### 3.2.1 How Uninterpreted Functions Are Used

Replacing functions with uninterpreted functions in a given formula is a common technique for making it easier to reason about (e.g., to prove its validity).

<sup>2</sup> From here on, we refer only to uninterpreted functions. Uninterpreted predicates are treated in a similar way.

At the same time, this process makes the formula *weaker*, which means that it can make a valid formula invalid. This observation is summarized in the following relation, where  $\varphi^{\text{UF}}$  is derived from a formula  $\varphi$  by replacing some or all of its functions with uninterpreted functions:

$$\models \varphi^{\text{UF}} \implies \models \varphi. \quad (3.3)$$

Uninterpreted functions are widely used in calculus and other branches of mathematics, but in the context of reasoning and verification, they are mainly used for simplifying proofs. Under certain conditions, uninterpreted functions let us reason about systems while ignoring the semantics of some or all functions, assuming they are not necessary for the proof. What does it mean to ignore the semantics of a function? (A formal explanation is briefly given in the aside on p. 61.) One way to look at this question is through the axioms that the function can be defined by. Ignoring the semantics of the function means that an interpretation neednot satisfy these axioms in order to satisfy the formula. The only thing it needs to satisfy is an axiom stating that the uninterpreted function, like any function, is *consistent*, i.e., given the same inputs, it returns the same outputs. This is the requirement of **functional consistency** (also called **functional congruence**):

**Functional consistency:** Instances of the same function return the same value if given equal arguments.

There are many cases in which the formula of interest is valid regardless of the interpretation of a function. In these cases, uninterpreted functions simplify the proof significantly, especially when it comes to mechanical proofs with the aid of automatic theorem provers.

Assume that we have a method for checking the validity of an EUF formula. Relying on this assumption, the basic scheme for using uninterpreted functions is the following:

1. Let  $\varphi$  denote a formula of interest that has interpreted functions. Assume that a validity check of  $\varphi$  is too hard (computationally), or even impossible.
2. Assign an uninterpreted function to each interpreted function in  $\varphi$ . Substitute each function in  $\varphi$  with the uninterpreted function to which it is mapped. Denote the new formula by  $\varphi^{\text{UF}}$ .
3. Check the validity of  $\varphi^{\text{UF}}$ . If it is valid, return “ $\varphi$  is valid” (this is justified by (3.3)). Otherwise, return “don’t know”.

The transformation in step 2 comes at a price, of course, as it loses information. As mentioned earlier, it causes the procedure to be incomplete, even if the original formula belongs to a decidable logic. When there exists a decision procedure for the input formula but it is too computationally hard to solve, one can design a procedure in which uninterpreted functions are gradually substituted back to their interpreted versions. We shall discuss this option further in Sect. 3.4.

$\varphi^{\text{UF}}$

### 3.2.2 An Example: Proving Equivalence of Programs

As a motivating example, consider the problem of proving the equivalence of the two C functions shown in Fig. 3.1. More specifically, the goal is to prove that they return the same value for every possible input `in`.

<pre>int power3(int in) { int i, out_a;   out_a = in;   for (i = 0; i &lt; 2; i++)     out_a = out_a * in;   return out_a; }</pre>	<pre>int power3_new(int in) { int out_b;    out_b = (in * in) * in;   return out_b; }</pre>
(a)	(b)

**Fig. 3.1.** Two C functions. The proof of their equivalence is simplified by replacing the multiplications (“\*”) in both programs with uninterpreted functions

In general, proving the equivalence of two programs is undecidable, which means that there is no sound and complete method to prove such an equivalence. In the present case, however, equivalence can be decided.<sup>3</sup> A key observation about these programs is that they have only bounded loops, and therefore it is possible to compute their input/output relations. The derivation of these relations from these two programs can be done as follows:

1. Remove the variable declarations and “return” statements.
2. Unroll the `for` loop.
3. Replace the left-hand side variable in each assignment with a new auxiliary variable.
4. Wherever a variable is read (referred to in an expression), replace it with the auxiliary variable that replaced it in the last place where it was assigned.
5. Conjoin all program statements.

These operations result in the two formulas  $\varphi_a$  and  $\varphi_b$ , which are shown in Fig. 3.2.<sup>4</sup>

It is left to show that these two I/O relations are actually equivalent, that is, to prove the validity of

$$\varphi_a \wedge \varphi_b \implies out2\_a = out0\_b . \quad (3.4)$$

<sup>3</sup> The undecidability of program verification and program equivalence is caused by unbounded memory usage, which does not occur in this example.

<sup>4</sup> A generalization of this form of translation to programs with “if” branches and other constructs is known as **static-single-assignment**(SSA). SSA is used in most optimizing compilers and can be applied to the verification of programs with bounded loops in popular programming languages such as C (see [107]). See also Example 1.25.

$$\begin{array}{l|l}
\begin{array}{l}
out0\_a = in \quad \wedge \\
out1\_a = out0\_a * in \wedge \\
out2\_a = out1\_a * in \\
\\
(\varphi_a)
\end{array}
&
\begin{array}{l}
out0\_b = (in*in)*in; \\
\\
(\varphi_b)
\end{array}
\end{array}$$

**Fig. 3.2.** Two formulas corresponding to the programs (a) and (b) in Fig. 3.1. The variables are defined over finite-width integers (i.e., bit vectors)

Uninterpreted functions can help in proving the equivalence of the programs (a) and (b), following the general scheme suggested in Sect. 3.2.1. The motivation in this case is computational: deciding formulas with multiplication over, for example, 32-bit variables is notoriously hard. Replacing the multiplication symbol with uninterpreted functions can solve the problem.

$$\begin{array}{l|l}
\begin{array}{l}
out0\_a = in \quad \wedge \\
out1\_a = G(out0\_a, in) \wedge \\
out2\_a = G(out1\_a, in) \\
\\
(\varphi_a^{UF})
\end{array}
&
\begin{array}{l}
out0\_b = G(G(in, in), in) \\
\\
(\varphi_b^{UF})
\end{array}
\end{array}$$

**Fig. 3.3.** After replacing “\*” with the uninterpreted function  $G$

Figure 3.3 presents  $\varphi_a^{UF}$  and  $\varphi_b^{UF}$ , which are  $\varphi_a$  and  $\varphi_b$  after the multiplication function has been replaced with a new uninterpreted function  $G$ . Similarly, if we also had addition, we could replace all of its instances with another uninterpreted function, say  $F$ . Instead of validating (3.4), we can now attempt to validate

$$\varphi_a^{UF} \wedge \varphi_b^{UF} \implies out2\_a = out0\_b. \quad (3.5)$$

Alternative methods to prove the equivalence of these two programs are discussed in the aside on p. 65. Other examples of the use of uninterpreted functions are presented in Sect. 3.5.

### 3.3 From Uninterpreted Functions to Equality Logic

Luckily, we do not need to examine all possible interpretations of an uninterpreted function in a given EUF formula in order to know whether it is valid. Instead, we rely on the strongest property that is common to all functions, namely functional consistency.<sup>5</sup> Relying on this property, we can reduce the decision problem of EUF formulas to that of deciding equality logic. We shall

<sup>5</sup> Note that the term *function* here refers to the mathematical definition. The situation is more complicated when considering functions in programming languages

**Aside: Alternative Decision Procedures**

The procedure in Sect. 3.2.2 is not the only way to automatically prove the equivalence of programs (a) and (b), of course. In this case, substitution is sufficient: by simply substituting  $out2\_a$  by  $out1\_a * in$ ,  $out1\_a$  by  $out0\_a * in$ , and  $out0\_a$  by  $in$  in  $\varphi_a$ , we can quickly (and automatically) prove (3.4), as we obtain syntactically equal expressions. However, there are many cases where such substitution is not efficient, as it can increase the size of the formula exponentially. It is also possible that substitution alone may be insufficient to prove equivalence. Consider, for example, the two functions `power3_con` and `power3_con_new`:

```

int power3_con                int power3_con_new
  (int in, int con) {         (int in, int con) {
  int i, out_a;               int out_b;
  out_a = in;
  for (i = 0; i < 2; i++)
    out_a = con?out_a * in    out_b = con?(in*in)*in
                        :out_a;           :in;
  return out_a;              return out_b;
}                               }
      (a)                       (b)

```

After substitution, we obtain two expressions,

$$out\_a = con? ((con? in * in : in) * in) : (con? in * in : in) \quad (3.6)$$

and

$$out\_b = con? (in * in) * in : in , \quad (3.7)$$

corresponding to the two functions. Not only are these two expressions not syntactically equivalent, but also the first expression grows exponentially with the number of iterations.

Another possible way to prove equivalence is to rely on the fact that the loops in the above programs are finite, and that the variables, as in any C program, are of finite type (e.g., integers are typically represented using 32-bit bit vectors – see Chap. 6). Therefore, the set of states reachable by the two programs can be represented and searched. This method can almost never compete, however, with decision procedures for equality logic and uninterpreted functions in terms of efficiency. There is a tradeoff, then, between efficiency and completeness.

see two possible reductions, **Ackermann's reduction** and **Bryant's reduction**, both of which enforce functional consistency. The former is somewhat more intuitive to understand, but also imposes certain restrictions on the decision procedures that can be used to solve it, unlike the latter. The implications of the differences between the two methods are explained in Sect. 4.6.

In the discussion that follows, for the sake of simplicity, we make several assumptions regarding the input formula: it has a single uninterpreted function, with a single argument, and no two instances of this function have the same argument. The generalization of the reductions is rather straightforward, as the examples later on demonstrate.

### 3.3.1 Ackermann's Reduction

Ackermann's reduction (Algorithm 3.3.1) adds explicit constraints to the formula in order to enforce the functional consistency requirement stated above. The algorithm reads an EUF formula  $\varphi^{\text{UF}}$  that we wish to validate, and transforms it to an equality logic formula  $\varphi^{\text{E}}$  of the form

$$\varphi^{\text{E}} := FC^{\text{E}} \implies flat^{\text{E}}, \quad (3.8)$$

where  $FC^{\text{E}}$  is a conjunction of functional-consistency constraints, and  $flat^{\text{E}}$  is a flattening of  $\varphi^{\text{UF}}$ , i.e., a formula in which each unique function instance is replaced with a corresponding new variable.

**Example 3.4.** Consider the formula

$$(x_1 \neq x_2) \vee (F(x_1) = F(x_2)) \vee (F(x_1) \neq F(x_3)), \quad (3.9)$$

which we wish to reduce to equality logic using Algorithm 3.3.1.

After assigning indices to the instances of  $F$  (for this example, we assume that this is done from left to right), we compute  $flat^{\text{E}}$  and  $FC^{\text{E}}$  accordingly:

$$flat^{\text{E}} := (x_1 \neq x_2) \vee (f_1 = f_2) \vee (f_1 \neq f_3), \quad (3.10)$$

$$\begin{aligned} FC^{\text{E}} := & (x_1 = x_2 \implies f_1 = f_2) \wedge \\ & (x_1 = x_3 \implies f_1 = f_3) \wedge \\ & (x_2 = x_3 \implies f_2 = f_3) . \end{aligned} \quad (3.11)$$

Equation (3.9) is valid if and only if the resulting equality formula is valid:

$$\varphi^{\text{E}} := FC^{\text{E}} \implies flat^{\text{E}}. \quad (3.12)$$

▀

---

such as C or JAVA. Functional consistency is guaranteed in that case only if we consider all the data that the function may read (including global variables, static variables, and data read from the environment) as argument of the function, and provided that the program is single-threaded.

**Algorithm 3.3.1: ACKERMANN'S-REDUCTION**

**Input:** An EUF formula  $\varphi^{\text{UF}}$  with  $m$  instances of an uninterpreted function  $F$

**Output:** An equality logic formula  $\varphi^{\text{E}}$  such that  $\varphi^{\text{E}}$  is valid if and only if  $\varphi^{\text{UF}}$  is valid

1. Assign indices to the uninterpreted-function instances from subexpressions outwards. Denote by  $F_i$  the instance of  $F$  that is given the index  $i$ , and by  $\text{arg}(F_i)$  its single argument.
2. Let  $\text{flat}^{\text{E}} \doteq \mathcal{T}(\varphi^{\text{UF}})$ , where  $\mathcal{T}$  is a function that takes an EUF formula (or term) as input and transforms it to an equality formula (or term, respectively) by replacing each uninterpreted-function instance  $F_i$  with a new term-variable  $f_i$  (in the case of nested functions, only the variable corresponding to the most external instance remains).
3. Let  $FC^{\text{E}}$  denote the following conjunction of functional consistency constraints:

$$FC^{\text{E}} := \bigwedge_{i=1}^{m-1} \bigwedge_{j=i+1}^m (\mathcal{T}(\text{arg}(F_i)) = \mathcal{T}(\text{arg}(F_j))) \implies f_i = f_j .$$

4. Let

$$\varphi^{\text{E}} := FC^{\text{E}} \implies \text{flat}^{\text{E}} .$$

Return  $\varphi^{\text{E}}$ .

 $m$  $F_i$  $\text{arg}(F_i)$  $\text{flat}^{\text{E}}$  $\mathcal{T}$  $FC^{\text{E}}$ 

In the next example, we go back to our running example for this chapter, and transform it to equality logic.

**Example 3.5.** Recall our main example. We left it in Fig. 3.3 after adding the uninterpreted-function symbol  $G$ . Now, using Ackermann's reduction, we can reduce it into an equality logic formula. This example also demonstrates how to generalize the reduction to functions with several arguments: only if all arguments of a pair of function instances are the same (pairwise), the return value of the function is forced to be the same.

Our example has four instances of the uninterpreted function  $G$ ,

$$G(\text{out0}_a, \text{in}), G(\text{out1}_a, \text{in}), G(\text{in}, \text{in}), \text{ and } G(G(\text{in}, \text{in}), \text{in}),$$

which we number in this order. On the basis of (3.5), we compute  $\text{flat}^{\text{E}}$ , replacing each uninterpreted-function symbol with the corresponding variable:

$$\text{flat}^{\text{E}} := \left( \left( \begin{array}{l} \text{out0}_a = \text{in} \wedge \\ \text{out1}_a = g_1 \wedge \\ \text{out2}_a = g_2 \end{array} \right) \wedge \text{out0}_b = g_4 \right) \implies \text{out2}_a = \text{out0}_b .$$

(3.13)



**Aside: Checking the *Satisfiability* of  $\varphi^{\text{UF}}$** 

Ackermann's reduction was defined above for checking the *validity* of  $\varphi^{\text{UF}}$ . It tells us that we need to check for the validity of  $\varphi^{\text{E}} := FC^{\text{E}} \implies flat^{\text{E}}$  or, equivalently, check that  $\neg\varphi^{\text{E}} := FC^{\text{E}} \wedge \neg flat^{\text{E}}$  is unsatisfiable. This is important in our case, because all the algorithms that we shall see later check for *satisfiability* of formulas, not for their validity. Thus, as a first step we need to negate  $\varphi^{\text{E}}$ .

What if we want to check for the satisfiability of  $\varphi^{\text{UF}}$ ? The short answer is that we need to check for the satisfiability of

$$\varphi^{\text{E}} := FC^{\text{E}} \wedge flat^{\text{E}}.$$

This is interesting. Normally, if we check for the satisfiability or validity of a formula, this corresponds to checking for the satisfiability of the formula or of its negation, respectively. Thus, we could expect that checking the satisfiability of  $\varphi^{\text{UF}}$  is equivalent to checking satisfiability of  $(FC^{\text{E}} \implies flat^{\text{E}})$ . However, this is not the same as the above equation. So what has happened here? The reason for the difference is that we check the satisfiability of  $\varphi^{\text{UF}}$  *before* the reduction. This means that we can use Ackermann's reduction to check the validity of  $\neg\varphi^{\text{UF}}$ . The functional-consistency constraints  $FC^{\text{E}}$  remain unchanged whether we check  $\varphi^{\text{UF}}$  or its negation  $\neg\varphi^{\text{UF}}$ . Thus, we need to check the validity of  $FC^{\text{E}} \implies \neg flat^{\text{E}}$ , which is the same as checking the satisfiability of  $FC^{\text{E}} \wedge flat^{\text{E}}$ , as stated above.

The functional-consistency constraints are given by

$$\begin{aligned} FC^{\text{E}} := & ((out0\_a = out1\_a \wedge in = in) \implies g_1 = g_2) \wedge \\ & ((out0\_a = in \wedge in = in) \implies g_1 = g_3) \wedge \\ & ((out0\_a = g_3 \wedge in = in) \implies g_1 = g_4) \wedge \\ & ((out1\_a = in \wedge in = in) \implies g_2 = g_3) \wedge \\ & ((out1\_a = g_3 \wedge in = in) \implies g_2 = g_4) \wedge \\ & ((in = g_3 \wedge in = in) \implies g_3 = g_4). \end{aligned} \tag{3.14}$$

The resulting equality formula is  $FC^{\text{E}} \implies flat^{\text{E}}$ , which we need to validate.

The reader may observe that most of these constraints are in fact redundant. The validity of the formula depends on  $G(out0\_a, in)$  being equal to  $G(in, in)$ , and  $G(out1\_a, in)$  being equal to  $G(G(in, in), in)$ . Hence, only the second and fifth constraints in (3.14) are necessary. In practice, such observations are important because the quadratic growth in the number of functional-consistency constraints may become a bottleneck. When comparing two systems, as in this case, it is frequently possible to detect in polynomial time large sets of constraints that can be removed without affecting the validity of the formula. More details of this technique can be found in [156].  $\blacksquare$

Finally, we consider the case in which there is more than one function symbol.

**Example 3.6.** Consider now the following formula, which we wish to validate:

$$x_1 = x_2 \implies \underbrace{F(\underbrace{F(\underbrace{G(x_1)})}_{f_1})}_{f_2} = \underbrace{F(\underbrace{F(\underbrace{G(x_2)})}_{f_3})}_{f_4}. \quad (3.15)$$

We index the function instances from the inside out (from subexpressions outwards) and compute the following:

$$flat^E := x_1 = x_2 \implies f_2 = f_4 \quad (3.16)$$

$$\begin{aligned} FC^E := x_1 = x_2 &\implies g_1 = g_2 \wedge \\ g_1 = f_1 &\implies f_1 = f_2 \wedge \\ g_1 = g_2 &\implies f_1 = f_3 \wedge \\ g_1 = f_3 &\implies f_1 = f_4 \wedge \\ f_1 = g_2 &\implies f_2 = f_3 \wedge \\ f_1 = f_3 &\implies f_2 = f_4 \wedge \\ g_2 = f_3 &\implies f_3 = f_4. \end{aligned} \quad (3.17)$$

Then, again,

$$\varphi^E := FC^E \implies flat^E. \quad (3.18)$$

▀

From these examples, it is clear how to generalize Algorithm 3.3.1 to multiple uninterpreted functions. We leave this and other extensions as an exercise (Problem 3.3).

### 3.3.2 Bryant's Reduction

*Bryant's reduction* (Algorithm 3.3.2) has the same goal as Ackermann's reduction: to transform EUF formulas to equality logic formulas, such that both are equivalent. To check the *satisfiability* of  $\varphi^{UF}$  rather than the validity, we return  $FC^E \wedge flat^E$  in the last step.

The semantics of the *case* expression used in step 3 is such that its value is determined by the first condition that is evaluated to TRUE. Its translation to an equality logic formula, assuming that the argument of  $F_i$  is a variable  $x_i$  for all  $i$ , is given by

$$\bigvee_{j=1}^i (F_i^* = f_j \wedge (x_j = x_i) \wedge \bigwedge_{k=1}^{j-1} (x_k \neq x_i)). \quad (3.22)$$

**Example 3.7.** Given the case expression

$$F_3^* = \left( \begin{array}{l} \text{case } x_1 = x_3 : f_1 \\ \quad \quad \quad x_2 = x_3 : f_2 \\ \quad \quad \quad \text{TRUE} : f_3 \end{array} \right), \quad (3.23)$$



simplifying the writing of the formula. We can do without them if we remove  $FC^E$  from the formula altogether and substitute them in  $flat^E$  with their definitions. The reason that we maintain them is to make the presentation more readable and to maintain a structure similar to Ackermann's reduction.

3. The definition of  $FC^E$ , which enforces functional consistency, relies on *case* expressions rather than on a pairwise enforcing of consistency.

The generalization of Algorithm 3.3.2 to functions with multiple arguments is straightforward, as we shall soon see in the examples.

**Example 3.8.** Let us return to our main example of this chapter, the problem of proving the equivalence of programs (a) and (b) in Fig. 3.1. We continue from Fig. 3.3, where the logical formulas corresponding to these programs are given, with the use of the uninterpreted function  $G$ . On the basis of (3.5), we compute  $flat^E$ , replacing each uninterpreted-function symbol with the corresponding variable:

$$flat^E := \left( \left( \begin{array}{l} out0\_a = in \wedge \\ out1\_a = G_1^* \wedge \\ out2\_a = G_2^* \end{array} \right) \wedge (out0\_b = G_4^*) \right) \implies out2\_a = out0\_b . \quad (3.25)$$

Not surprisingly, this looks very similar to (3.13). The only difference is that instead of the  $g_i$  variables, we now have the  $G_i^*$  macros, for  $1 \leq i \leq 4$ . Recall their origin: the function instances are  $G(out0\_a, in)$ ,  $G(out1\_a, in)$ ,  $G(in, in)$  and  $G(G(in, in), in)$ , which we number in this order. The corresponding functional-consistency constraints are

$$FC^E := \left( \begin{array}{l} G_1^* = g_1 \\ G_2^* = \left( \begin{array}{ll} \text{case } out0\_a = out1\_a \wedge in = in : g_1 & \\ \text{TRUE} & : g_2 \end{array} \right) \\ G_3^* = \left( \begin{array}{ll} \text{case } out0\_a = in \wedge in = in : g_1 & \\ out1\_a = in \wedge in = in : g_2 & \\ \text{TRUE} & : g_3 \end{array} \right) \\ G_4^* = \left( \begin{array}{ll} \text{case } out0\_a = G_3^* \wedge in = in : g_1 & \\ out1\_a = G_3^* \wedge in = in : g_2 & \\ in = G_3^* \wedge in = in : g_3 & \\ \text{TRUE} & : g_4 \end{array} \right) \end{array} \right) \wedge \quad (3.26)$$

and since we are checking for validity, the formula to be checked is

$$\varphi^E := FC^E \implies flat^E . \quad (3.27)$$

▀

**Example 3.9.** If there are multiple uninterpreted-function symbols, the reduction is applied to each of them separately, as demonstrated in the following example, in which we consider the formula of Example 3.6 again:

$$x_1 = x_2 \implies \underbrace{\underbrace{F(F(G(x_1)))}_{F_1^*}}_{F_2^*} = \underbrace{\underbrace{F(F(G(x_2)))}_{F_3^*}}_{F_4^*}. \quad (3.28)$$

As before, we number the function instances of each of the uninterpreted-function symbols  $F$  and  $G$  from the inside out (this order is required in Bryant's reduction). Applying Bryant's reduction, we obtain

$$flat^E := (x_1 = x_2 \implies F_2^* = F_4^*), \quad (3.29)$$

$$\begin{aligned} FC^E := & F_1^* = f_1 \quad \wedge \\ & F_2^* = \left( \begin{array}{ll} \text{case } G_1^* = F_1^* : f_1 & \\ \text{TRUE} & : f_2 \end{array} \right) \wedge \\ & F_3^* = \left( \begin{array}{ll} \text{case } G_1^* = G_2^* : f_1 & \\ F_1^* = G_2^* : f_2 & \\ \text{TRUE} & : f_3 \end{array} \right) \wedge \\ & F_4^* = \left( \begin{array}{ll} \text{case } G_1^* = F_3^* : f_1 & \\ F_1^* = F_3^* : f_2 & \\ G_2^* = F_3^* : f_3 & \\ \text{TRUE} & : f_4 \end{array} \right) \wedge \\ & G_1^* = g_1 \quad \wedge \\ & G_2^* = \left( \begin{array}{ll} \text{case } x_1 = x_2 : g_1 & \\ \text{TRUE} & : g_2 \end{array} \right), \end{aligned} \quad (3.30)$$

and

$$\varphi^E := FC^E \implies flat^E. \quad (3.31)$$

Note that in any satisfying assignment that satisfies  $x_1 = x_2$  (the premise of (3.28)),  $F_1^*$  and  $F_3^*$  are equal to  $f_1$ , while  $F_2^*$  and  $F_4^*$  are equal to  $f_2$ .  $\blacksquare$

The difference between Ackermann's and Bryant's reductions is not just syntactic, as was hinted earlier. It has implications for the decision procedure that one can use when solving the resulting formula. We discuss this point further in Sect. 4.6.

### 3.4 Functional Consistency Is Not Enough

Functional consistency is not always sufficient for proving correct statements. This is not surprising, as we clearly lose information by replacing concrete, interpreted functions with uninterpreted functions. Consider, for example, the *plus* ('+') function. Now suppose that we are given a formula containing the two function instances  $x_1 + y_1$  and  $x_2 + y_2$ , and, owing to other parts of the formula, it holds that  $x_1 = y_2$  and  $y_1 = x_2$ . Further, suppose that we

replace “+” with a binary uninterpreted function  $F$ . Since in Algorithms 3.3.1 and 3.3.2 we only compare arguments pairwise in the order in which they appear, the proof cannot rely on the fact that these two function instances are evaluated to give the same result. In other words, the functional-consistency constraints alone do not capture the commutativity of the “+” function, which may be necessary for the proof. This demonstrates the fact that by using uninterpreted functions we lose completeness (see Definition 1.6).

One may add, of course, additional constraints that capture more information about the original function – commutativity, in the case of the example above. For example, considering Ackermann’s reduction for the above example, let  $f_1, f_2$  be the variables that encode the two function instances, respectively. We can then replace the functional-consistency constraint for this pair with the stronger constraint

$$((x_1 = x_2 \wedge y_1 = y_2) \vee (x_1 = y_2 \wedge y_1 = x_2)) \implies f_1 = f_2 . \quad (3.32)$$

Such constraints can be tailored as needed, to reflect properties of the uninterpreted functions. In other words, by adding these constraints we make them **partially interpreted functions**, as we model some of their properties. For the multiplication function, for example, we can add a constraint that if one of the arguments is equal to 0, then so is the result. Generally, the more abstract the formula is, the easier it is, computationally, to solve it. On the other hand, the more abstract the formula is, the fewer correct facts about its original version can be proven. The right abstraction level for a given formula can be found by a trial-and-error process. Such a process can even be automated with an **abstraction-refinement loop**,<sup>6</sup> as can be seen in Algorithm 3.4.1 (this is not so much an algorithm as a framework that needs to be concretized according to the exact problem at hand). In step 2, the algorithm returns “Valid” if the abstract formula is valid. The correctness of this step is implied by (3.3). If, on the other hand, the formula is not valid and the abstract formula  $\varphi'$  is identical to the original one, the algorithm returns “Valid” in the next step. The optional step that follows (step 4) is not necessary for the soundness of the algorithm, but only for its performance. This step is worth executing only if it is easier than solving  $\varphi$  itself.

Plenty of room for creativity is left when one is implementing such an algorithm: which constraints to add in step 5? When to resort to the original interpreted functions? How to implement step 4? An instance of such a procedure is described, for the case of bit-vector arithmetic, in Sect. 6.3.

---

<sup>6</sup> Abstraction-refinement loops [111] are implemented in many *model checkers* [46] (tools for verifying temporal properties of transition systems) and other automated formal-reasoning tools. The types of abstractions used can be very different than from those presented here, but the basic elements of the iterative process are the same.

**Aside: Rewriting systems**

Observations such as “a multiplication by 0 is equal to 0” can be formulated with *rewriting rules*. Such rules are the basis of *rewriting systems* [64, 99], which are used in several branches of mathematics and mathematical logic. Rewriting systems, in their basic form, define a set of terms and (possibly non-deterministic) rules for transforming them. Theorem provers that are based on rewriting systems (such as ACL2 [104]) use hundreds of such rules. Many of these rules can be used in the context of the partially interpreted functions that were studied in Sect. 3.4, as demonstrated for the “multiply by 0” rule.

Rewriting systems, as a formalism, have the same power as a Turing machine. They are frequently used for defining and implementing inference systems, for simplifying formulas by replacing subexpressions with equal but simpler subexpressions, for computing results of arithmetic expressions, and so forth. Such implementations require the design of a strategy for applying the rules, and a mechanism based on pattern matching for detecting the set of applicable rules at each step.

**Algorithm 3.4.1: ABSTRACTION-REFINEMENT**

**Input:** A formula  $\varphi$  in a logic  $L$ , such that there is a decision procedure for  $L$  with uninterpreted functions

**Output:** “Valid” if  $\varphi$  is valid, “Not valid” otherwise

1.  $\varphi' := \mathcal{T}(\varphi)$ .
2. If  $\varphi'$  is valid then return “Valid”.
3. If  $\varphi' = \varphi$  then return “Not valid”.
4. (Optional) Let  $\alpha'$  be a counterexample to the validity of  $\varphi'$ . If it is possible to derive a counterexample  $\alpha$  to the validity of  $\varphi$  (possibly by extending  $\alpha'$  to those variables in  $\varphi$  that are not in  $\varphi'$ ), return “Not valid”.
5. Refine  $\varphi'$  by adding more constraints as discussed in sect. 3.4, or by replacing uninterpreted functions with their original interpreted versions (reaching, in the worst case, the original formula  $\varphi$ ).
6. Return to step 2.

**3.5 Two Examples of the Use of Uninterpreted Functions**

Uninterpreted functions can be used for *property-based* verification, that is, proving that a certain property holds for a given model. Occasionally it happens that properties are correct regardless of the semantics of a certain function, and functional consistency is all that is needed for the proof. In such cases, replacing the function with an uninterpreted function can simplify the proof.

The more common use of uninterpreted functions, however, is for proving *equivalence* between systems. In the chip design industry, proving equivalence between two versions of a hardware circuit is a standard procedure. Another application is **translation validation**, a process of proving the semantic equivalence of the input and output of a compiler. Indeed, we end this chapter with a detailed description of these two problem domains.

In both applications, it is expected that every function on one side of the equation can be mapped to a similar function on the other side. In such cases, replacing all functions with an uninterpreted version and using one of the reductions that we saw in Sects. 3.3.1 and 3.3.2 is typically sufficient for proving equivalence.

### 3.5.1 Proving Equivalence of Circuits

*Pipelining* is a technique for improving the performance of a circuit such as a microprocessor. The computation is split into phases, called pipeline stages. This allows one to speed up the computation by making use of concurrent computation, as is done in an assembly line in a factory.

The clock frequency of a circuit is limited by the length of the longest path between latches (i.e., memory components), which is, in the case of a pipelined circuit, simply the length of the longest stage. The delay of each path is affected by the gates along that path and the delay that each one of them imposes.

Figure 3.4(a) shows a pipelined circuit. The input, denoted by  $in$ , is processed in the first stage. We model the combinational gates within the stages with uninterpreted functions, denoted by  $C, F, G, H, K$ , and  $D$ . For the sake of simplicity, we assume that they each impose the same delay. The circuit applies function  $F$  to the inputs  $in$ , and stores the result in latch  $L_1$ . This can be formalized as follows:

$$L_1 = F(in) . \quad (3.33)$$

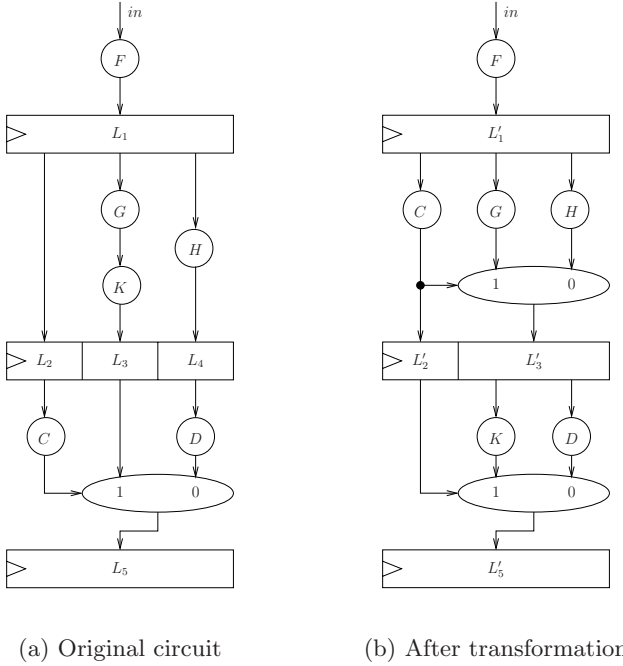
The second stage computes values for  $L_2, L_3$ , and  $L_4$ :

$$\begin{aligned} L_2 &= L_1 , \\ L_3 &= K(G(L_1)) , \\ L_4 &= H(L_1) . \end{aligned} \quad (3.34)$$

The third stage contains a *multiplexer*. A multiplexer is a circuit that selects between two inputs according to the value of a Boolean signal. In this case, this selection signal is computed by a function  $C$ . The output of the multiplexer is stored in latch  $L_5$ :

$$L_5 = C(L_2) ? L_3 : D(L_4) . \quad (3.35)$$





**Fig. 3.4.** Showing the correctness of a transformation of a pipelined circuit using uninterpreted functions. After the transformation, the circuit has a shorter longest path between stages, and thus can be operated at a higher clock frequency

Observe that the second stage contains two functions,  $G$  and  $K$ , where the output of  $G$  is used as an input for  $K$ . Suppose that this is the longest path within the circuit. We now aim to transform the circuit in order to make it work faster. This can be done in this case by moving the gates represented by  $K$  down into the third stage.

Observe also that only one of the values in  $L_3$  and  $L_4$  is used, as the multiplexer selects one of them depending on  $C$ . We can therefore remove one of the latches by introducing a second multiplexer in the second stage. The circuit after these changes is shown in Fig. 3.4(b). It can be formalized as follows:

$$\begin{aligned}
 L'_1 &= F(in) , \\
 L'_2 &= C(L'_1) , \\
 L'_3 &= C(L'_1) ? G(L'_1) : H(L'_1) , \\
 L'_5 &= L'_2 ? K(L'_3) : D(L'_3) .
 \end{aligned}
 \tag{3.36}$$

The final result of the computation is stored in  $L_5$  in the original circuit, and in  $L'_5$  in the modified circuit. We can show that the transformations are correct by proving that for all inputs, the conjunction of the above equalities implies

$$L_5 = L'_5 . \quad (3.37)$$

This proof can be automated by using a decision procedure for equalities and uninterpreted functions.

### 3.5.2 Verifying a Compilation Process with Translation Validation

The next example illustrates a translation validation process that relies on uninterpreted functions and Ackermann's reduction. Unlike the hardware example, we start from interpreted functions and replace them with uninterpreted functions.

Suppose that a source program contains the statement

$$z = (x_1 + y_1) * (x_2 + y_2) , \quad (3.38)$$

which the compiler that we wish to check compiles into the following sequence of three assignments:

$$u_1 = x_1 + y_1; u_2 = x_2 + y_2; z = u_1 * u_2 . \quad (3.39)$$

Note the two new auxiliary variables  $u_1$  and  $u_2$  that have been added by the compiler. To verify this translation, we construct the verification condition

$$u_1 = x_1 + y_1 \wedge u_2 = x_2 + y_2 \wedge z = u_1 * u_2 \implies z = (x_1 + y_1) * (x_2 + y_2) , \quad (3.40)$$

whose validity we wish to check.<sup>7</sup>

We now abstract the concrete functions appearing in the formula, namely addition and multiplication, by the abstract uninterpreted-function symbols  $F$  and  $G$ , respectively. The abstracted version of the implication above is

$$\begin{aligned} (u_1 = F(x_1, y_1) \wedge u_2 = F(x_2, y_2) \wedge z = G(u_1, u_2)) \\ \implies z = G(F(x_1, y_1), F(x_2, y_2)) . \end{aligned} \quad (3.41)$$

Clearly, if the abstracted version is valid, then so is the original concrete one (see (3.3)).

Next, we apply Ackermann's reduction (Algorithm 3.3.1), replacing each function by a new variable, but adding, for each pair of terms with the same function symbol, an extra antecedent that guarantees the functionality of these terms. Namely, if the two arguments of the original terms are equal, then the terms should be equal.

<sup>7</sup> This verification condition is an implication rather than an equivalence because we are attempting to prove that the values allowed in the target code are also allowed in the source code, but not necessarily the other way. This asymmetry can be relevant when the source code is interpreted as a specification that allows multiple behaviors, only one of which is actually implemented. For the purpose of demonstrating the use of uninterpreted functions, whether we use an implication or an equivalence is immaterial.

Applying Ackermann's reduction to the abstracted formula, we obtain the following equality formula:

$$\varphi^E := \left( \begin{array}{l} (x_1 = x_2 \wedge y_1 = y_2 \implies f_1 = f_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \implies g_1 = g_2) \end{array} \right) \implies \left( \begin{array}{l} (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1) \implies \\ z = g_2 \end{array} \right), \quad (3.42)$$

which we can rewrite as

$$\varphi^E := \left( \begin{array}{l} (x_1 = x_2 \wedge y_1 = y_2 \implies f_1 = f_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \implies g_1 = g_2) \wedge \\ u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \end{array} \right) \implies z = g_2. \quad (3.43)$$

It is left to prove, then, the validity of this equality logic formula.

The success of such a process depends on how different the two sides are. Suppose that we are attempting to perform translation validation for a compiler that does not perform heavy arithmetic optimizations. In such a case, the scheme above will probably succeed. If, on the other hand, we are comparing two *arbitrary* source codes, even if they are equivalent, it is unlikely that the same scheme will be sufficient. It is possible, for example, that one side uses the function  $2 * x$  while the other uses  $x + x$ . Since addition and multiplication are represented by two different uninterpreted functions, they are not associated with each other in any way according to Algorithm 3.3.1, and hence the proof of equivalence is not able to rely on the fact that the two expressions are semantically equal.

## 3.6 Problems

### 3.6.1 Warm-up Exercises

**Problem 3.1 (practicing Ackermann's and Bryant's reductions).**

Given the formula

$$\begin{array}{l} F(F(x_1)) \neq F(x_1) \wedge \\ F(F(x_1)) \neq F(x_2) \wedge \\ x_2 = F(x_1), \end{array} \quad (3.44)$$

reduce its validity problem to a validity problem of an equality logic formula through Ackermann's reduction and Bryant's reduction.

### 3.6.2 Problems

**Problem 3.2 (eliminating constants).** Prove that given an equality logic formula, Algorithm 3.1.1 returns an equisatisfiable formula without constants.<sup>8</sup>

<sup>8</sup> Further discussion of the constants-elimination problem appears in the next chapter, as part of Problem 4.4.

**Problem 3.3 (Ackermann's reduction).** Extend Algorithm 3.3.1 to multiple function symbols and to functions with multiple arguments.

**Problem 3.4 (Bryant's reduction).** Suppose that in Algorithm 3.3.2, the definition of  $F_i$  is replaced by

$$F_i^* = \left( \begin{array}{l} \text{case } \mathcal{T}^*(\text{arg}(F_1^*)) = \mathcal{T}^*(\text{arg}(F_i^*)) : F_1^* \\ \vdots \\ \mathcal{T}^*(\text{arg}(F_{i-1}^*)) = \mathcal{T}^*(\text{arg}(F_i^*)) : F_{i-1}^* \\ \text{TRUE} : f_i \end{array} \right), \quad (3.45)$$

the difference being that the terms on the right refer to the  $F_j^*$  variables,  $1 \leq j < i$ , rather than to the  $f_j$  variables. Does this change the value of  $F_i^*$ ? Prove a negative answer or give an example.

**Problem 3.5 (abstraction/refinement).** Frequently, the functional-consistency constraints become the bottleneck in the verification procedure, as their number is quadratic in the number of function instances. In such cases, even solving the first iteration of Algorithm 3.4.1 is too hard.

Show an abstraction/refinement algorithm that begins with  $flat^E$  and gradually adds functional-consistency constraints.

*Hint:* note that given an assignment  $\alpha'$  that satisfies a formula with only some of the functional-consistency constraints, checking whether  $\alpha'$  respects functional consistency is not trivial. This is because  $\alpha'$  does not necessarily refer to all variables (if the formula contains nested functions, some may disappear in the process of abstraction). Hence  $\alpha'$  cannot be tested directly against a version of the formula that contains all functional-consistency constraints.

## 3.7 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$\varphi^E$	Equality formula	60
$C_c$	A variable used for substituting a constant $c$ in the process of removing constants from equality formulas	60
$\varphi^{UF}$	Equality formula + uninterpreted functions (before reduction to equality logic)	62
<i>continued on next page</i>		

<i>continued from previous page</i>		
<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$\mathcal{T}$	A function that transforms an input formula or term by replacing each uninterpreted function $F_i$ with a new variable $f_i$	67
$FC^E$	Functional-consistency constraints	67
$\mathcal{T}^*$	A function similar to $\mathcal{T}$ , that replaces each uninterpreted function $F_i$ with $F_i^*$	70
$flat^E$	Equal to $\mathcal{T}(\varphi^{UF})$ in Ackermann's reduction, and to $\mathcal{T}^*(\varphi^{UF})$ in Bryant's reduction	67, 70
$F_i^*$	In Bryant's reduction, a macro variable representing the case expression associated with the function instance $F_i()$ that was substituted by $F_i$	70

## Decision Procedures for Equality Logic and Uninterpreted Functions

In Chap. 3, we saw how useful the theory of equality logic with uninterpreted-function (EUF) is. In this chapter, we concentrate on decision procedures for EUF and on algorithms for simplifying EUF formulas. Recall that we are solving the satisfiability problem for formulas in negation normal form (NNF – see Definition 1.10) without constants, as those can be removed with, for example, Algorithm 3.1.1. With the exception of Sect. 4.1, we handle equality logic without uninterpreted functions, assuming that these are eliminated by one of the reduction methods introduced in Chap. 3.

### 4.1 Deciding a Conjunction of Equalities and Uninterpreted Functions with Congruence Closure

We begin by showing a method for solving a conjunction of equalities and uninterpreted functions. As is the case for most of the theories that we consider in this book, the satisfiability problem for conjunctions of predicates can be solved in polynomial time.

Note that a decision procedure for a conjunction of equality predicates is not sufficient to support uninterpreted functions as well, as both Ackermann's and Bryant's reductions (Chap. 3) introduce disjunctions into the formula.

As an alternative, Shostak proposed in 1978 a method for handling uninterpreted functions directly. Starting from a conjunction  $\varphi^{\text{EUF}}$  of equalities and disequalities over variables and uninterpreted functions, he proposed a two-stage algorithm (see Algorithm 4.1.1), which is based on computing equivalence classes. The version of the algorithm that is presented here assumes that the uninterpreted functions have a single argument. The extension to the general case is left as an exercise (Problem 4.3).

**Algorithm 4.1.1:** CONGRUENCE-CLOSURE

**Input:** A conjunction  $\varphi^{\text{UF}}$  of equality predicates over variables and uninterpreted functions

**Output:** “Satisfiable” if  $\varphi^{\text{UF}}$  is satisfiable, and “Unsatisfiable” otherwise

1. Build congruence-closed equivalence classes.
  - (a) Initially, put two terms  $t_1, t_2$  (either variables or uninterpreted-function instances) in their own equivalence class if  $(t_1 = t_2)$  is a predicate in  $\varphi^{\text{UF}}$ . All other variables form singleton equivalence classes.
  - (b) Given two equivalence classes with a shared term, merge them. Repeat until there are no more classes to be merged.
  - (c) Compute the *congruence closure*: given two terms  $t_i, t_j$  that are in the same class and that  $F(t_i)$  and  $F(t_j)$  are terms in  $\varphi^{\text{UF}}$  for some uninterpreted function  $F$ , merge the classes of  $F(t_i)$  and  $F(t_j)$ . Repeat until there are no more such instances.
2. If there exists a disequality  $t_i \neq t_j$  in  $\varphi^{\text{UF}}$  such that  $t_i$  and  $t_j$  are in the same equivalence class, return “Unsatisfiable”. Otherwise return “Satisfiable”.

**Example 4.1.** Consider the conjunction

$$\varphi^{\text{UF}} := x_1 = x_2 \wedge x_2 = x_3 \wedge x_4 = x_5 \wedge x_5 \neq x_1 \wedge F(x_1) \neq F(x_3). \quad (4.1)$$

Initially, the equivalence classes are

$$\{x_1, x_2\}, \{x_2, x_3\}, \{x_4, x_5\}, \{F(x_1)\}, \{F(x_3)\}. \quad (4.2)$$

Step 1(b) of Algorithm 4.1.1 merges the first two classes:

$$\{x_1, x_2, x_3\}, \{x_4, x_5\}, \{F(x_1)\}, \{F(x_3)\}. \quad (4.3)$$

The next step also merges the classes containing  $F(x_1)$  and  $F(x_3)$ , because  $x_1$  and  $x_2$  are in the same class:

$$\{x_1, x_2, x_3\}, \{x_4, x_5\}, \{F(x_1), F(x_3)\}. \quad (4.4)$$

In step 2, we note that  $F(x_1) \neq F(x_3)$  is a predicate in  $\varphi^{\text{UF}}$ , but that  $F(x_1)$  and  $F(x_3)$  are in the same class. Hence,  $\varphi^{\text{UF}}$  is unsatisfiable.  $\blacksquare$

Variants of Algorithm 4.1.1 can be implemented efficiently with a **union-find** data structure, which results in a time complexity of  $O(n \log n)$  (see, for example, [141]).

In the original presentation of his method, Shostak implemented support for disjunctions by means of case-splitting, which is the bottleneck in this method. For example, given the formula

$$\varphi^{\text{UF}} := x_1 = x_2 \vee (x_2 = x_3 \wedge x_4 = x_5 \wedge x_5 \neq x_1 \wedge F(x_1) \neq F(x_3)), \quad (4.5)$$

he considered separately the two cases corresponding to the left and right parts of the disjunction. This can work well as long as there are not too many cases to consider.

The more interesting question is how to solve the general case efficiently, where the given formula has an arbitrary Boolean structure. This problem arises with all the theories that we study in this book. There are two main approaches. A highly efficient method is to combine a SAT solver with an algorithm such as Algorithm 4.1.1, where the former searches for a satisfying assignment to the Boolean *skeleton* of the formula (an abstraction of the formula where each unique predicate is replaced with a new Boolean variable), and the latter is used to check whether this assignment corresponds to a satisfying assignment to the equality predicates – we dedicate Chap. 11 to this technique. A second approach is based on a full reduction to propositional logic, and is the subject of the rest of this chapter.

## 4.2 Basic Concepts

In this section, we present several basic terms that are used later in the chapter. We assume from here on that uninterpreted functions have already been eliminated, i.e., that we are solving the satisfiability problem for equality logic without uninterpreted functions. Recall that we are also assuming that the formula is given to us in NNF and without constants. Recall further that an atom in such formulas is an equality predicate, and a literal is either an atom or its negation (see Definition 1.11). Given an equality logic formula  $\varphi^{\text{E}}$ , we denote the set of atoms of  $\varphi^{\text{E}}$  by  $At(\varphi^{\text{E}})$ .

**Definition 4.2 (equality and disequality literals sets).** *The equality literals set  $E_ =$  of an equality logic formula  $\varphi^{\text{E}}$  is the set of positive literals in  $\varphi^{\text{E}}$ . The disequality literals set  $E_{\neq}$  of an equality logic formula  $\varphi^{\text{E}}$  is the set of disequality literals in  $\varphi^{\text{E}}$ .*

It is possible, of course, that an equality may appear in the equality literals set and its negation in the disequality literals set.

**Example 4.3.** Consider the negation normal form of  $\neg\varphi^{\text{E}}$  in (3.43):

$$\neg\varphi^{\text{E}} := \left( \begin{array}{l} (x_1 \neq x_2 \vee y_1 \neq y_2 \vee f_1 = f_2) \wedge \\ (u_1 \neq f_1 \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1) \end{array} \right) \wedge z \neq g_2. \quad (4.6)$$

At( $\varphi^{\text{E}}$ )

$E_ =$

$E_{\neq}$



We therefore have

$$\begin{aligned} E_=&:= \{(f_1 = f_2), (g_1 = g_2), (u_1 = f_1), (u_2 = f_2), (z = g_1)\} \\ E_{\neq}&:= \{(x_1 \neq x_2), (y_1 \neq y_2), (u_1 \neq f_1), (u_2 \neq f_2), (z \neq g_2)\}. \end{aligned} \quad (4.7)$$


▀

$G^E$

**Definition 4.4 (equality graph).** Given an equality logic formula  $\varphi^E$  in NNF, the equality graph that corresponds to  $\varphi^E$ , denoted by  $G^E(\varphi^E)$ , is an undirected graph  $(V, E_=: E_=: E_{\neq})$  where the nodes in  $V$  correspond to the variables in  $\varphi^E$ , the edges in  $E_=:$  correspond to the predicates in the equality literals set of  $\varphi^E$  and the edges in  $E_{\neq}$  correspond to the predicates in the disequality literals set of  $\varphi^E$ .

Note that we overload the symbols  $E_=:$  and  $E_{\neq}$  so that each represents both the literals sets and the edges that represent them in the equality graph. Similarly, when we say that an assignment “satisfies an edge”, we mean that it satisfies the literal represented by that edge.

We may write simply  $G^E$  for an equality graph when the formula it corresponds to is clear from the context. Graphically, equality literals are represented as dashed edges and disequality literals as solid edges, as illustrated in Fig. 4.1.



**Fig. 4.1.** An equality graph. Dashed edges represent  $E_=:$  literals (equalities), and solid edges represent  $E_{\neq}$  literals (disequalities)

It is important to note that the equality graph  $G^E(\varphi^E)$  represents an *abstraction* of  $\varphi^E$ : more specifically, it represents all the equality logic formulas that have the same literals sets as  $\varphi^E$ . Since it disregards the Boolean connectives, it can represent both a satisfiable and an unsatisfiable formula. For example, although  $x_1 = x_2 \wedge x_1 \neq x_2$  is unsatisfiable and  $x_1 = x_2 \vee x_1 \neq x_2$  is satisfiable, both formulas are represented by the same equality graph.

**Definition 4.5 (equality path).** An equality path in an equality graph  $G^E$  is a path consisting of  $E_=:$  edges. We denote by  $x =^* y$  the fact that there exists an equality path from  $x$  to  $y$  in  $G^E$ , where  $x, y \in V$ .

**Definition 4.6 (disequality path).** A disequality path in an equality graph  $G^E$  is a path consisting of  $E_=:$  edges and a single  $E_{\neq}$  edge. We denote by  $x \neq^* y$  the fact that there exists a disequality path from  $x$  to  $y$  in  $G^E$ , where  $x, y \in V$ .

$x =^* y$

$x \neq^* y$

Similarly, we use the terms *simple equality path* and *simple disequality path* when the path is required to be loop-free.

Consider Fig. 4.1 and observe, for example, that  $x_2 =^* x_4$  owing to the path  $x_2, x_5, x_4$ , and  $x_2 \neq^* x_4$  owing to the path  $x_2, x_5, x_1, x_4$ . In this case, both paths are simple. Intuitively, if  $x =^* y$  in  $G^E(\varphi^E)$ , then it might be necessary to assign the two variables equal values in order to satisfy  $\varphi^E$ . We say “might” because, once again, the equality graph obscures details about  $\varphi^E$ , as it disregards the Boolean structure of  $\varphi^E$ . The only fact that we know from  $x =^* y$  is that there exist formulas whose equality graph is  $G^E(\varphi^E)$  and that in any assignment satisfying them,  $x = y$ . However, we do not know whether  $\varphi^E$  is one of them. A disequality path  $x \neq^* y$  in  $G^E(\varphi^E)$  implies the opposite: it might be necessary to assign different values to  $x$  and  $y$  in order to satisfy  $\varphi^E$ .

The case in which both  $x =^* y$  and  $x \neq^* y$  hold in  $G^E(\varphi^E)$  requires special attention. We say that the graph, in this case, contains a *contradictory cycle*.

**Definition 4.7 (contradictory cycle).** *In an equality graph, a contradictory cycle is a cycle with exactly one disequality edge.*

For every pair of nodes  $x, y$  in a contradictory cycle, it holds that  $x =^* y$  and  $x \neq^* y$ .

Contradictory cycles are of special interest to us because the conjunction of the literals corresponding to their edges is unsatisfiable. Furthermore, since we have assumed that there are no constants in the formula, these are the only topologies that have this property. Consider, for example, a contradictory cycle with nodes  $x_1, \dots, x_k$  in which  $(x_1, x_k)$  is the disequality edge. The conjunction

$$x_1 = x_2 \wedge \dots \wedge x_{k-1} = x_k \wedge x_k \neq x_1 \quad (4.8)$$

is clearly unsatisfiable.

All the decision procedures that we consider refer explicitly or implicitly to contradictory cycles. For most algorithms we can further simplify this definition by considering only *simple contradictory cycles*. A cycle is simple if it is represented by a path in which none of the vertices is repeated, other than the starting and ending vertices.

### 4.3 Simplifications of the Formula


Regardless of the algorithm that is used for deciding the satisfiability of a given equality logic formula  $\varphi^E$ , it is almost always the case that  $\varphi^E$  can be simplified a great deal before the algorithm is invoked. Algorithm 4.3.1 presents such a simplification.

**Algorithm 4.3.1:** SIMPLIFY-EQUALITY-FORMULA

- Input:** An equality formula  $\varphi^E$   
**Output:** An equality formula  $\varphi^{E'}$  equisatisfiable with  $\varphi^E$ , with length less than or equal to the length of  $\varphi^E$
1. Let  $\varphi^{E'} := \varphi^E$ .
  2. Construct the equality graph  $G^E(\varphi^{E'})$ .
  3. Replace each pure literal in  $\varphi^{E'}$  whose corresponding edge is not part of a simple contradictory cycle with TRUE.
  4. Simplify  $\varphi^{E'}$  with respect to the Boolean constants TRUE and FALSE (e.g., replace  $\text{TRUE} \vee \phi$  with TRUE, and  $\text{FALSE} \wedge \phi$  with FALSE).
  5. If any rewriting has occurred in the previous two steps, go to step 2.
  6. Return  $\varphi^{E'}$ .

The following example illustrates the steps of Algorithm 4.3.1.

**Example 4.8.** Consider (4.6). Figure 4.2 illustrates  $G^E(\varphi^E)$ , the equality graph corresponding to  $\varphi^E$ .



**Fig. 4.2.** The equality graph corresponding to Example 4.8. The edges  $f_1 = f_2$ ,  $x_1 \neq x_2$  and  $y_1 \neq y_2$  are not part of any contradictory cycle, and hence their respective predicates in the formula can be replaced with TRUE

In this case, the edges  $f_1 = f_2$ ,  $x_1 \neq x_2$  and  $y_1 \neq y_2$  are not part of any simple contradictory cycle and can therefore be substituted by TRUE. This results in

$$\varphi^{E'} := \left( \begin{array}{l} (\text{TRUE} \vee \text{TRUE} \vee \text{TRUE}) \wedge \\ (u_1 \neq f_1 \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \wedge z \neq g_2) \end{array} \right), \quad (4.9)$$

which, after simplification according to step 4, is equal to

$$\varphi^{E'} := \left( \begin{array}{l} (u_1 \neq f_1 \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \wedge z \neq g_2) \end{array} \right). \quad (4.10)$$

Reconstructing the equality graph after this simplification does not yield any more simplifications, and the algorithm terminates.

Now, consider a similar formula in which the predicates  $x_1 \neq x_2$  and  $u_1 \neq f_1$  are swapped. This results in the formula

$$\varphi^E := \left( \begin{array}{l} (u_1 \neq f_1 \vee y_1 \neq y_2 \vee f_1 = f_2) \wedge \\ (x_1 \neq x_2 \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \wedge z \neq g_2) \end{array} \right). \quad (4.11)$$


Although we start from exactly the same graph, the simplification algorithm is now much more effective. After the first step we have

$$\varphi^{E'} := \left( \begin{array}{l} (u_1 \neq f_1 \vee \text{TRUE} \vee \text{TRUE}) \wedge \\ (\text{TRUE} \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \wedge z \neq g_2) \end{array} \right), \quad (4.12)$$

which, after step 4, simplifies to

$$\varphi^{E'} := \left( (u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \wedge z \neq g_2) \right). \quad (4.13)$$

The graph corresponding to  $\varphi^{E'}$  after this step appears in Fig. 4.3.



**Fig. 4.3.** An equality graph corresponding to (4.13), showing the first iteration of step 4

Clearly, no edges in  $\varphi^{E'}$  belong to a contradictory cycle after this step, which implies that we can replace all the remaining predicates by TRUE. Hence, in this case, simplification alone proves that the formula is satisfiable, without invoking a decision procedure.  $\blacksquare$

Although we leave the formal proof of the correctness of Algorithm 4.3.1 as an exercise (Problem 4.5), let us now consider what such a proof may look like. Correctness can be shown by proving that steps 3 and 4 maintain satisfiability (as these are the only steps in which the formula is changed). The simplifications in step 4 trivially maintain satisfiability, so the main problem is step 3.

Let  $\varphi_1^E$  and  $\varphi_2^E$  be the equality formulas before and after step 3, respectively. We need to show that these formulas are equisatisfiable.

( $\Rightarrow$ ) If  $\varphi_1^E$  is satisfiable, then so is  $\varphi_2^E$ . This is implied by the monotonicity of NNF formulas (see Theorem 1.14) and the fact that only pure literals are replaced by TRUE.

( $\Leftarrow$ ) If  $\varphi_2^E$  is satisfiable, then so is  $\varphi_1^E$ . Only a proof sketch and an example will be given here. The idea is to construct a satisfying assignment  $\alpha_1$  for  $\varphi_1^E$  while relying on the existence of a satisfying assignment  $\alpha_2$  for  $\varphi_2^E$ . Specifically,  $\alpha_1$  should satisfy exactly the same predicates as are satisfied by  $\alpha_2$ , but also satisfy all those predicates that were replaced by TRUE. The following simple observation can be helpful in this construction: given a satisfying assignment to an equality formula, shifting the values in the assignment uniformly maintains satisfaction (because the values of the equality predicates remain the same). The same observation applies to an assignment of *some* of the variables, as long as none of the predicates that refer to *one* of these variables becomes FALSE owing to the new assignment.

Consider, for example, (4.11) and (4.12), which correspond to  $\varphi_1^E$  and  $\varphi_2^E$ , respectively, in our argument. An example of a satisfying assignment to the latter is

$$\alpha_2 := \{u_1 \mapsto 0, f_1 \mapsto 0, f_2 \mapsto 1, u_2 \mapsto 1, z \mapsto 0, g_1 \mapsto 0, g_2 \mapsto 1\}. \quad (4.14)$$

First,  $\alpha_1$  is set equal to  $\alpha_2$ . Second, we need to extend  $\alpha_1$  with an assignment of those variables not assigned by  $\alpha_2$ . The variables in this category are  $x_1, x_2, y_1$ , and  $y_2$ , which can be trivially satisfied because they are not part of any equality predicate. Hence, assigning a unique value to each of them is sufficient. For example, we can now have

$$\alpha_1 := \alpha_1 \cup \{x_1 \mapsto 2, x_2 \mapsto 3, y_1 \mapsto 4, y_2 \mapsto 5\}. \quad (4.15)$$

Third, we need to consider predicates that are replaced by TRUE in step 3 but are not satisfied by  $\alpha_1$ . In our example,  $f_1 = f_2$  is such a predicate. To solve this problem, we simply shift the assignment to  $f_2$  and  $u_2$  so that the predicate  $f_1 = f_2$  is satisfied (a shift by minus 1 in this case). This clearly maintains the satisfaction of the predicate  $u_2 = f_2$ . The assignment that satisfies  $\varphi_1^E$  is thus

$$\alpha_1 := \{u_1 \mapsto 0, f_1 \mapsto 0, f_2 \mapsto 0, u_2 \mapsto 0, z \mapsto 0, g_1 \mapsto 0, g_2 \mapsto 1, \\ x_1 \mapsto 2, x_2 \mapsto 3, y_1 \mapsto 4, y_2 \mapsto 5\}. \quad (4.16)$$

A formal proof based on this argument should include a precise definition of these shifts, i.e., which vertices do they apply to, and an argument as to why no circularity can occur. Circularity can affect the termination of the procedure that constructs  $\alpha_1$ .

## 4.4 A Graph-Based Reduction to Propositional Logic

We now consider a decision procedure for equality logic that is based on a reduction to propositional logic. This procedure was originally presented by Bryant and Velev in [39] (under the name of the **sparse method**). Several definitions and observations are necessary.

**Definition 4.9 (nonpolar equality graph).** Given an equality logic formula  $\varphi^E$ , the nonpolar equality graph corresponding to  $\varphi^E$ , denoted by  $G_{\text{NP}}^E(\varphi^E)$ , is an undirected graph  $(V, E)$  where the nodes in  $V$  correspond to the variables in  $\varphi^E$ , and the edges in  $E$  correspond to  $\text{At}(\varphi^E)$ , i.e., the equality predicates in  $\varphi^E$ .

 $G_{\text{NP}}^E$ 

A nonpolar equality graph represents a degenerate version of an equality graph (Definition 4.4), since it disregards the polarity of the equality predicates.

Given an equality logic formula  $\varphi^E$ , the procedure generates two propositional formulas  $e(\varphi^E)$  and  $\mathcal{B}_{\text{trans}}$ , such that

 $e(\varphi^E)$ 

$$\varphi^E \text{ is satisfiable} \iff e(\varphi^E) \wedge \mathcal{B}_{\text{trans}} \text{ is satisfiable.} \quad (4.17)$$

 $\mathcal{B}_{\text{trans}}$ 

The formulas  $e(\varphi^E)$  and  $\mathcal{B}_{\text{trans}}$  are defined as follows:

- The formula  $e(\varphi^E)$  is the **propositional skeleton** of  $\varphi^E$ , which means that every equality predicate of the form  $x_i = x_j$  in  $\varphi^E$  is replaced with a new Boolean variable  $e_{i,j}$ .<sup>1</sup> For example, let

$$\varphi^E := x_1 = x_2 \wedge (((x_2 = x_3) \wedge (x_1 \neq x_3)) \vee (x_1 \neq x_2)). \quad (4.18)$$

Then,

$$e(\varphi^E) := e_{1,2} \wedge ((e_{2,3} \wedge \neg e_{1,3}) \vee \neg e_{1,2}). \quad (4.19)$$

It is not hard to see that if  $\varphi^E$  is satisfiable, then so is  $e(\varphi^E)$ . The other direction, however, does not hold. For example, while (4.18) is unsatisfiable, its encoding in (4.19) is satisfiable. To maintain an equisatisfiability relation, we need to add constraints that impose the transitivity of equality, which was lost in the encoding. This is the role of  $\mathcal{B}_{\text{trans}}$ .

- The formula  $\mathcal{B}_{\text{trans}}$  is a conjunction of implications, which are called *transitivity constraints*. Each such implication is associated with a cycle in the nonpolar equality graph. For a cycle with  $n$  edges,  $\mathcal{B}_{\text{trans}}$  forbids an assignment FALSE to one of the edges when all the other edges are assigned TRUE. Imposing this constraint for each of the edges in each one of the cycles is sufficient to satisfy the condition stated in (4.17).

**Example 4.10.** The atoms  $x_1 = x_2, x_2 = x_3, x_1 = x_3$  form a cycle of size 3 in the nonpolar equality graph. The following constraint is sufficient for maintaining the condition stated in (4.17):

$$\mathcal{B}_{\text{trans}} = \left( \begin{array}{l} (e_{1,2} \wedge e_{2,3} \implies e_{1,3}) \wedge \\ (e_{1,2} \wedge e_{1,3} \implies e_{2,3}) \wedge \\ (e_{2,3} \wedge e_{1,3} \implies e_{1,2}) \end{array} \right). \quad (4.20)$$

▀

<sup>1</sup> To avoid introducing dual variables such as  $e_{i,j}$  and  $e_{j,i}$ , we can assume that all equality predicates in  $\varphi^E$  appear in such a way that the left variable precedes the right one in some predefined order.

Adding  $n$  constraints for each cycle is not very practical, however, because there can be an exponential number of cycles in a given undirected graph.

**Definition 4.11 (chord).** *A chord of a cycle is an edge connecting two non-adjacent nodes of the cycle. If a cycle has no chords in a given graph, it is called a chord-free cycle.*

Bryant and Velev proved the following theorem:

**Theorem 4.12.** *It is sufficient to add transitivity constraints over simple chord-free cycles in order to maintain (4.17).*

For a formal proof, see [39]. The following example may be helpful for developing an intuition as to why this theorem is correct.

**Example 4.13.** Consider the cycle  $(x_3, x_4, x_8, x_7)$  in one of the two graphs in Fig. 4.4. It contains the chord  $(x_3, x_8)$  and, hence, is not chord-free. Now assume that we wish to assign TRUE to all edges in this cycle other than  $(x_3, x_4)$ . If  $(x_3, x_8)$  is assigned TRUE, then the assignment to the simple chord-free cycle  $(x_3, x_4, x_8)$  contradicts transitivity. If  $(x_3, x_8)$  is assigned FALSE, then the assignment to the simple chord-free cycle  $(x_3, x_7, x_8)$  contradicts transitivity. Thus, the constraints over the chord-free cycles are sufficient for preventing the transitivity-violating assignment to the cycle that includes a chord. ▀


The number of simple chord-free cycles in a graph can still be exponential in the number of vertices. Hence, building  $\mathcal{B}_{trans}$  such that it directly constrains every such cycle can make the size of this formula exponential in the number of variables. Luckily, we have:

**Definition 4.14 (chordal graphs).** *A chordal graph is an undirected graph in which no cycle of size 4 or more is chord-free.*

Every graph can be made chordal in a time polynomial in the number of vertices.<sup>2</sup> Since the only chord-free cycles in a chordal graph are triangles, this implies that applying Theorem 4.12 to such a graph results in a formula of size not more than cubic in the number of variables (three constraints for each triangle in the graph). The newly added chords are represented by new variables that appear in  $\mathcal{B}_{trans}$  but not in  $e(\varphi^E)$ . Algorithm 4.4.1 summarizes the steps of this method.

**Example 4.15.** Figure 4.4 depicts a nonpolar equality graph before and after making it chordal. We use solid edges, but note that these *should not be confused with the solid edges in (polar) equality graphs*, where they denote

<sup>2</sup> We simply remove all vertices from the graph one by one, each time connecting the neighbors of the eliminated vertex if they were not already connected. The original graph plus the edges added in this process is a chordal graph.



**Fig. 4.4.** A nonchordal nonpolar equality graph corresponding to  $\varphi^E$  (left), and a possible chordal version of it (right)

**Algorithm 4.4.1:** EQUALITY-LOGIC-TO-PROPOSITIONAL-LOGIC

**Input:** An equality formula  $\varphi^E$

**Output:** A propositional formula equisatisfiable with  $\varphi^E$

1. Construct a Boolean formula  $e(\varphi^E)$  by replacing each atom of the form  $x_i = x_j$  in  $\varphi^E$  with a Boolean variable  $e_{i,j}$ .
2. Construct the nonpolar equality graph  $G_{\text{NP}}^E(\varphi^E)$ .
3. Make  $G_{\text{NP}}^E(\varphi^E)$  chordal.
4.  $\mathcal{B}_{\text{trans}} := \text{TRUE}$ .
5. For each triangle  $(e_{i,j}, e_{j,k}, e_{i,k})$  in  $G_{\text{NP}}^E(\varphi^E)$ ,

$$\begin{aligned} \mathcal{B}_{\text{trans}} &:= \mathcal{B}_{\text{trans}} \wedge \\ &\quad (e_{i,j} \wedge e_{j,k} \implies e_{i,k}) \wedge \\ &\quad (e_{i,j} \wedge e_{i,k} \implies e_{j,k}) \wedge \\ &\quad (e_{i,k} \wedge e_{j,k} \implies e_{i,j}) . \end{aligned} \tag{4.21}$$

6. Return  $e(\varphi^E) \wedge \mathcal{B}_{\text{trans}}$ .

disequalities. After the graph has been made chordal, it contains four triangles and, hence,  $\mathcal{B}_{\text{trans}}$  conjoins 12 constraints. For example, for the triangle  $(x_1, x_2, x_5)$ , the constraints are

$$\begin{aligned} e_{1,2} \wedge e_{2,5} &\implies e_{1,5} , \\ e_{1,5} \wedge e_{2,5} &\implies e_{1,2} , \\ e_{1,2} \wedge e_{1,5} &\implies e_{2,5} . \end{aligned} \tag{4.22}$$

The added edge  $e_{2,5}$  corresponds to a new auxiliary variable  $e_{2,5}$  that appears in  $\mathcal{B}_{\text{trans}}$  but not in  $e(\varphi^E)$ .  $\blacksquare$

There exists a version of this algorithm that is based on the (polar) equality graph, and generates a smaller number of transitivity constraints. See Problem 4.6 for more details.



## 4.5 Equalities and Small-Domain Instantiations

In this section, we show a method for solving equality logic formulas by relying on the **small-model property** that this logic has. This means that every satisfiable formula in this logic has a model (a satisfying interpretation) of finite size. Furthermore, in equality logic there is a computable bound on the size of such a model. We use the following definitions in the rest of the discussion.

**Definition 4.16 (adequacy of a domain for a formula).** *A domain is adequate for a formula if the formula either is unsatisfiable or has a model within this domain.*

**Definition 4.17 (adequacy of a domain for a set of formulas).** *A domain is adequate for a set of formulas if it is adequate for each formula in the set.*

In the case of equality logic, each set of formulas with the same number of variables has an easily computable adequate finite domain, as we shall soon see. The existence of such a domain immediately suggests a decision procedure: simply enumerate all assignments within this domain and check whether one of them satisfies the formula. Our solution strategy, therefore, for checking whether a given equality formula  $\varphi^E$  is satisfiable, can be summarized as follows:

1. Determine, in polynomial time, a **domain allocation**

$$D : \text{var}(\varphi^E) \mapsto 2^{\mathbb{N}} \quad (4.23)$$

(where  $\text{var}(\varphi^E)$  denotes the set of variables of  $\varphi^E$ ), by mapping each variable  $x_i \in \text{var}(\varphi^E)$  into a finite set of integers  $D(x_i)$ , such that  $\varphi^E$  is satisfiable if and only if it is satisfiable within  $D$  (i.e., there exists a satisfying assignment in which each variable  $x_i$  is assigned an integer from  $D(x_i)$ ).

2. Encode each variable  $x_i$  as an enumerated type over its finite domain  $D(x_i)$ . Construct a propositional formula representing  $\varphi^E$  under this finite domain, and use either BDDs or SAT to check if this formula is satisfiable.

This strategy is called **small-domain instantiation**, since we instantiate the variables with a finite set of values from the domain computed, each time checking whether it satisfies the formula. The number of instantiations in the worst case is what we call the size of the **state space** spanned by a domain. The size of the state space of a domain  $D$ , denoted by  $|D|$  is equal to the product of the numbers of elements in the domains of the individual variables. Clearly, the success of this method depends on its ability to find domain allocations with small state spaces.

### 4.5.1 Some Simple Bounds

We now show several bounds on the number of elements in an adequate domain. Let  $\Phi_n$  be the (infinite) set of all equality logic formulas with  $n$  variables and without constants.  $\Phi_n$

**Theorem 4.18 (“folk theorem”).** *The uniform domain allocation  $\{1, \dots, n\}$  for all  $n$  variables is adequate for  $\Phi_n$ .*

*Proof.* Let  $\varphi^E \in \Phi_n$  be a satisfiable equality logic formula. Every satisfying assignment  $\alpha$  to  $\varphi^E$  reflects a partition of its variables into equivalence classes. That is, two variables are in the same equivalence class if and only if they are assigned the same value by  $\alpha$ . Since there are only equalities and disequalities in  $\varphi^E$ , every assignment which reflects the same equivalence classes satisfies exactly the same predicates as  $\alpha$ . Since all partitions into equivalence classes over  $n$  variables are possible in the domain  $1, \dots, n$ , this domain is adequate for  $\varphi^E$ . ▀

This bound, although not yet tight, implies that we can encode each variable in a  $\Phi_n$  formula with no more than  $\lceil \log n \rceil$  bits, and with a total of  $n \lceil \log n \rceil$  bits for the entire formula in the worst case. This is very encouraging, because it is already better than the worst-case complexity of Algorithm 4.4.1, which requires  $n \cdot (n - 1)/2$  bits (one bit per pair of variables) in the worst case.

#### Aside: The Complexity Gap

Why is there a complexity gap between domain allocation and the encoding method that we described in Sect. 4.4? Where is the wasted work in EQUALITY-LOGIC-TO-PROPOSITIONAL-LOGIC? Both algorithms merely partition the variables into classes of equal variables, but they do it in a different way. Instead of asking ‘which *subset* of  $\{v_1, \dots, v_n\}$  is each variable equal to?’, with the domain-allocation technique we ask instead ‘which *value* in the range  $\{1, \dots, n\}$  is each variable equal to?’. For each variable, rather than exploring the range of subsets of  $\{v_1, \dots, v_n\}$  to which it may be equal, we instead explore the range of values  $\{1, \dots, n\}$ . The former requires one bit per element in this set, or a total of  $n$  bits, while the latter requires only  $\log n$  bits.

The domain  $1, \dots, n$ , as suggested above, results in a state space of size  $n^n$ . We can do better if we do not insist on a uniform domain allocation, which allocates the same domain to all variables.

**Theorem 4.19.** *Assume for each formula  $\varphi^E \in \Phi_n$ ,  $\text{var}(\varphi^E) = \{x_1, \dots, x_n\}$ . The domain allocation  $D := \{x_i \mapsto \{1, \dots, i\} \mid 1 \leq i \leq n\}$  is adequate for  $\Phi_n$ .*

*Proof.* As argued in the proof of Theorem 4.18, every satisfying assignment  $\alpha$  to  $\varphi^E \in \Phi_n$  reflects a partition of the variables to equivalence classes. We construct an assignment  $\alpha'$  as follows.

For each equivalence class  $C$ :

- Let  $x_i$  be the variable with the lowest index in  $C$ .
- Assign  $i$  to all the variables in  $C$ .

Since all the other variables in  $C$  have indices higher than  $i$ ,  $i$  is in their domain, and hence this assignment is feasible. Since each variable appears in exactly one equivalence class, every class of variables is assigned a different value, which means that  $\alpha'$  satisfies the same equality predicates as  $\alpha$ . This implies that  $\alpha'$  satisfies  $\varphi^E$ .  $\blacksquare$

The adequate domain suggested in Theorem 4.19 has a smaller state space, of size  $n!$ . In fact, it is conjectured that  $n!$  is also a lower bound on the size of domain allocations adequate for this class of formulas.

Let us now consider the case in which the formula contains constants.

$\Phi_{n,k}$

**Theorem 4.20.** *Let  $\Phi_{n,k}$  be the set of equality logic formulas with  $n$  variables and  $k$  constants. Assume, without loss of generality, that the constants are  $c_1 < \dots < c_k$ . The domain allocation*

$$D := \{x_i \mapsto \{c_1, \dots, c_k, c_k + 1, \dots, c_k + i\} \mid 1 \leq i \leq n\} \quad (4.24)$$

*is adequate for  $\Phi_{n,k}$ .*

The proof is left as an exercise (Problem 4.8).

The adequate domain suggested in Theorem 4.20 results in a state space of size  $(k+n)!/k!$ . As stated in Sect. 3.1.3, constants can be eliminated by adding more variables and constraints ( $k$  variables in this case), but note that this would result in a larger state space.

The next few sections are dedicated to an algorithm that reduces the allocated domain further, based on an analysis of the equality graph associated with the input formula.

— — —

*Sects. 4.5.2, 4.5.3, and 4.5.4 cover advanced topics.*

## 4.5.2 Graph-Based Domain Allocation

The formula sets  $\Phi_n$  and  $\Phi_{n,k}$  utilize only a simple structural characteristic common to all of their members, namely, the number of variables and constants. As a result, they group together many formulas of radically different nature. It is not surprising that the best size of adequate domain allocation for the whole set is so high. By paying attention to additional structural similarities of formulas, we can form smaller sets of formulas and obtain much smaller adequate domain allocations.

As before, we assume that  $\varphi^E$  is given in negation normal form. Let  $e$  denote a set of equality literals and  $\Phi(e)$  the set of all equality logic formulas whose literals set is equal to  $e$ . Let  $E(\varphi^E)$  denote the set of  $\varphi^E$ 's literals. Thus,  $\Phi(E(\varphi^E))$  is the set of all equality logic formulas that have the same set of literals as  $\varphi^E$ . Obviously,  $\varphi^E \in \Phi(E(\varphi^E))$ . Note that  $\Phi(e)$  can include both satisfiable and unsatisfiable formulas. For example, let  $e$  be the set

$$\{x_1 = x_2, x_1 \neq x_2\}. \quad (4.25)$$

Then  $\Phi(e)$  includes both the satisfiable formula

$$x_1 = x_2 \vee x_1 \neq x_2 \quad (4.26)$$

and the unsatisfiable formula

$$x_1 = x_2 \wedge x_1 \neq x_2. \quad (4.27)$$

An adequate domain, recall, is concerned only with the *satisfiable* formulas that can be constructed from literals in the set. Thus, we should not worry about (4.27). We should, however, be able to satisfy (4.26), as well as formulas such as  $x_1 = x_2 \wedge (\text{TRUE} \vee x_1 \neq x_2)$  and  $x_1 \neq x_2 \wedge (\text{TRUE} \vee x_1 = x_2)$ . One adequate domain for the set  $\Phi(e)$  is

$$D := \{x_1 \mapsto \{0\}, x_2 \mapsto \{0, 1\}\}. \quad (4.28)$$

It is not hard to see that this domain is minimal, i.e., there is no adequate domain with a state space smaller than 2 for  $\Phi(e)$ .

How do we know, then, which subsets of the literals in  $E(\varphi^E)$  we need to be able to satisfy within the domain  $D$ , in order for  $D$  to be adequate for  $\Phi(E(\varphi^E))$ ? The answer is that we need only to be able to satisfy *consistent* subsets of literals, i.e., subsets for which the conjunction of literals in each of them is satisfiable.

A set  $e$  of equality literals is consistent if and only if it does not contain one of the following two patterns:

1. A chain of the form  $x_1 = x_2, x_2 = x_3, \dots, x_{r-1} = x_r$  together with the formula  $x_1 \neq x_r$ .
2. A chain of the form  $c_1 = x_2, x_2 = x_3, \dots, x_{r-1} = c_r$  where  $c_1$  and  $c_r$  represent different constants.

In the equality graph corresponding to  $e$ , the first pattern appears as a contradictory cycle (Definition 4.7) and the second as an equality path (Definition 4.5) between two constants.

To summarize, a domain allocation  $D$  is adequate for  $\Phi(E(\varphi^E))$  if every consistent subset  $e \subseteq E(\varphi^E)$  is satisfiable within  $D$ . Hence, finding an adequate domain for  $\Phi(E(\varphi^E))$  is reduced to the following problem:

Associate with each variable  $x_i$  a set of integers  $D(x_i)$  such that every consistent subset  $e \in E(\varphi^E)$  can be satisfied with an assignment from these sets.

$\Phi(e)$

$E(\varphi^E)$

We wish to find sets of this kind that are as small as possible, in polynomial time.

### 4.5.3 The Domain Allocation Algorithm

Let  $G^E(\varphi^E)$  be the equality graph (see Definition 4.4) corresponding to  $\varphi^E$ , defined by  $(V, E_=, E_\neq)$ . Let  $G^E_=_$  and  $G^E_\neq$  denote two subgraphs of  $G^E(\varphi^E)$ , defined by  $(V, E_=_)$  and  $(V, E_\neq)$ , respectively. As before, we use dashed edges to represent  $G^E_=_$  edges and solid edges to represent  $G^E_\neq$  edges. A vertex is called *mixed* if it is adjacent to edges in both  $G^E_=_$  and  $G^E_\neq$ .

On the basis of the definitions above, Algorithm 4.5.1 computes an economical domain allocation  $D$  for the variables in a given equality formula  $\varphi^E$ . The algorithm receives as input the equality graph  $G^E(\varphi^E)$ , and returns as output a domain which is adequate for the set  $\Phi(E(\varphi^E))$ . Since  $\varphi^E \in \Phi(E(\varphi^E))$ , this domain is adequate for  $\varphi^E$ .

We refer to the values that were added in steps I.A.2, I.C, II.A.1, and II.B as the *characteristic* values of these vertices. We write  $char(x_i) = u_i$  and  $char(x_k) = u_{C_-}$ . Note that every vertex is assigned a single characteristic value. Vertices that are assigned their characteristic values in steps I.C and II.A.1 are called *individually assigned vertices*, whereas the vertices assigned characteristic values in step II.B are called *communally assigned vertices*. We assume that new values are assigned in ascending order, so that  $char(x_i) < char(x_j)$  implies that  $x_i$  was assigned its characteristic value before  $x_j$ . Consequently, we require that all new values are larger than the largest constant  $C_{max}$ . This assumption is necessary only for simplifying the proof in later sections.

The description of the algorithm presented above leaves open the order in which vertices are chosen in step II.A.1. This order has a strong impact on the size of the resulting state space. Since the values given in this step are distributed on the graph  $G^E_=_$  in step II.A.2, we would like to keep this set as small as possible. Furthermore, we would like to partition the graph quickly, in order to limit this distribution. A rather simple, but effective heuristic for this purpose is to choose vertices according to a greedy criterion, where mixed vertices are chosen in descending order of their degree in  $G^E_\neq$ . We denote the set of vertices chosen in step II.A.1 by  $MV$ , to remind ourselves that they are mixed vertices.

**Example 4.21.** We wish to check whether (4.6), copied below, is satisfiable:

$$\neg\varphi^E := \left( \begin{array}{l} (x_1 \neq x_2 \vee y_1 \neq y_2 \vee f_1 = f_2) \wedge \\ (u_1 \neq f_1 \vee u_2 \neq f_2 \vee g_1 = g_2) \wedge \\ u_1 = f_1 \wedge u_2 = f_2 \wedge z = g_1 \end{array} \right) \wedge z \neq g_2. \quad (4.29)$$

The sets  $E_=_$  and  $E_\neq$  are:

$$\begin{aligned} E_=_ &:= \{(f_1 = f_2), (g_1 = g_2), (u_1 = f_1), (u_2 = f_2), (z = g_1)\} \\ E_\neq &:= \{(x_1 \neq x_2), (y_1 \neq y_2), (u_1 \neq f_1), (u_2 \neq f_2), (z \neq g_2)\}, \end{aligned} \quad (4.30)$$

**Algorithm 4.5.1: DOMAIN-ALLOCATION-FOR-EQUALITIES**

**Input:** An equality graph  $G^E$

**Output:** An adequate domain (in the form of a set of integers for each variable-vertex) for the set of formulas over literals that are represented by  $G^E$  edges

**I. Eliminating constants and preprocessing**

Initially,  $D(x_i) = \emptyset$  for all vertices  $x_i \in G^E$ .

A. For each constant-vertex  $c_i$  in  $G^E$ , do:

1. (Empty item, for the sake of symmetry with step II.A.)
2. Assign  $D(x_j) := D(x_j) \cup \{c_i\}$  for each vertex  $x_j$ , such that there is an equality path from  $c_i$  to  $x_j$  not through any other constant-vertex.
3. Remove  $c_i$  and its adjacent edges from the graph.

B. Remove all  $G^E_{\neq}$  edges that do not lie on a contradictory cycle.

C. For every singleton vertex (a vertex comprising a connected component by itself)  $x_i$ , add to  $D(x_i)$  a new value  $u_i$ . Remove  $x_i$  and its adjacent edges from the graph.

**II. Value allocation**


A. While there are mixed vertices in  $G^E$  do:

1. Choose a mixed vertex  $x_i$ . Add  $u_i$ , a new value, to  $D(x_i)$ .
2. Assign  $D(x_j) := D(x_j) \cup \{u_i\}$  for each vertex  $x_j$ , such that there is an equality path from  $x_i$  to  $x_j$ .
3. Remove  $x_i$  and its adjacent edges from the graph.

B. For each (remaining) connected  $G^E_{=}$  component  $C_{=}$ , add a common new value  $u_{C_{=}}$  to  $D(x_k)$ , for every  $x_k \in C_{=}$ .

Return  $D$ .

and the corresponding equality graph  $G^E(\neg\varphi^E)$  reappears in Fig. 4.5.



**Fig. 4.5.** The equality graph  $G^E(\neg\varphi^E)$

We refrain in this example from applying preprocessing, in order to make the demonstration of the algorithm more informative and interesting. This example results in a state space of size  $11^{11}$  if we use the domain  $\{1, \dots, n\}$  as suggested in Theorem 4.18, and a state space of size  $11!$  ( $\approx 4 \times 10^7$ ) if we use the domain suggested in Theorem 4.19. Applying Algorithm 4.5.1, on the other hand, results in an adequate domain spanning a state space of size 48, as can be seen in Fig. 4.6.

Step	$x_1$	$x_2$	$y_1$	$y_2$	$u_1$	$f_1$	$f_2$	$u_2$	$g_2$	$z$	$g_1$	Removed
I.B												edges ( $x_1 - x_2$ ), ( $y_1 - y_2$ )
I.C	0	1	2	3								$x_1, x_2, y_1, y_2$
II.A					4	4	4	4				$f_1$
II.A							5	5				$f_2$
II.A									6	6	6	$g_2$
II.B					7							
II.B								8				
II.B										9	9	
Final $D$ -sets	0	1	2	3	4, 7	4	4, 5	4, 5, 8	6	6, 9	6, 9	State space = 48

Fig. 4.6. Application of Algorithm 4.5.1 to (4.29)

Using a small improvement concerning the new values allocated in step II.A.1, this allocation can be reduced further, down to a domain of size 16. This improvement is the subject of Problem 4.12.

For demonstration purposes, consider a formula  $\varphi^E$  where  $g_1$  is replaced by the constant “3”. In this case the component  $(z, g_1, g_2)$  is handled as follows: in step I.A, “3” is added to  $D(g_2)$  and  $D(z)$ . The edge  $(z, g_2)$ , now no longer part of a contradictory cycle, is then removed in step I.B and a distinct new value is added to each of these variables in step I.C. ▀

Algorithm 4.5.1 is polynomial in the size of the input graph: steps I.A and II.A are iterated a number of times not more than the number of vertices in the graph; step I.B is iterated not more than the number of edges in  $G_{\neq}^E$ ; steps I.A.2, I.B, II.A.2 and II.B can be implemented with depth-first search (DFS).

#### 4.5.4 A Proof of Soundness

In this section, we argue for the soundness of Algorithm 4.5.1. We begin by describing a procedure which, given the allocation  $D$  produced by this algorithm

and a consistent subset  $e$ , assigns to each variable  $x_i \in G^E$  an integer value  $a_e(x_i) \in D(x_i)$ . We then continue by proving that this assignment satisfies the literals in  $e$ .

$a_e$


### An Assignment Procedure

Given a consistent subset of literals  $e$  and its corresponding equality graph  $G^E(e)$ , assign to each variable-vertex  $x_i \in G^E(e)$  a value  $a_e(x_i) \in D(x_i)$ , according to the following rules:

- R1** If  $x_i$  is connected by a (possibly empty)  $G^E(e)$ -path to an individually assigned vertex  $x_j$ , assign to  $x_i$  the minimal value of  $char(x_j)$  among such  $x_j$ 's.
- R2** Otherwise, assign to  $x_i$  its communally assigned value  $char(x_i)$ .

To see why all vertices are assigned a value by this procedure, observe that every vertex is allocated a characteristic value before it is removed. This can be an individual characteristic value allocated in steps I.C and II.A.1, or a communal value allocated in step II.B. Every vertex  $x_i$  that has an individual characteristic value can be assigned a value  $a_e(x_i)$  by **R1**, because it has at least the empty equality path leading to an individually allocated vertex, namely itself. All other vertices are allocated a communal value that makes them eligible for a value assignment by **R2**.

**Example 4.22.** Consider the  $D$ -sets in Fig. 4.6. Let us apply the above assignment procedure to a consistent subset  $e$  that contains all edges, excluding the two edges between  $u_1$  and  $f_1$ , the dashed edge between  $g_1$  and  $g_2$ , and the solid edge between  $f_2$  and  $u_2$  (see Fig. 4.7).



**Fig. 4.7.** The consistent set of edges  $e$  considered in Example 4.22 and the values assigned by the assignment procedure

The assignment is as follows:

- By **R1**,  $x_1, x_2, y_1$  and  $y_2$  are assigned the characteristic values “0”, “1”, “2”, and “3”, respectively, which they received in step I.C.
- By **R1**,  $f_1, f_2$  and  $u_2$  are assigned the value  $char(f_1) = “4”$ , because  $f_1$  was the first mixed vertex in the subgraph  $\{f_1, f_2, u_2\}$  that was removed in step II.A, and consequently it has the minimal characteristic value.



- By **R1**,  $g_2$  is assigned the value  $\text{char}(g_2) = \text{“6”}$ , which it received in step II.A.
- By **R2**,  $z$  and  $g_1$  are assigned the value  $\text{“9”}$ , which they received in step II.B.
- By **R2**,  $u_1$  is assigned the value  $\text{“7”}$ , which it received in step II.B.

▀

**Theorem 4.23.** *The assignment procedure is feasible (i.e., the value assigned to a node by the procedure belongs to its  $D$ -set).*

*Proof.* Consider first the two classes of vertices that are assigned a value by **R1**. The first class includes vertices that are removed in step I.C. These vertices have only one (empty)  $G_{\perp}^E(e)$ -path to themselves, and are therefore assigned the characteristic value that they received in that step. The second class includes vertices that have a (possibly empty)  $G_{\perp}^E(e)$ -path to a vertex from  $MV$ . Let  $x_i$  denote such a vertex, and let  $x_j$  be the vertex with the minimal characteristic value that  $x_i$  can reach on  $G_{\perp}^E(e)$ . Since  $x_i$  and all the vertices on this path were still part of the graph when  $x_j$  was removed in step II.A, then  $\text{char}(x_j)$  was added to  $D(x_i)$  according to step II.A.2. Thus, the assignment of  $\text{char}(x_j)$  to  $x_i$  is feasible.

Next, consider the vertices that are assigned a value by **R2**. Every vertex that was removed in step I.C or II.A is clearly assigned a value by **R1**. All the other vertices were communally assigned a value in step II.B. In particular, the vertices that do not have a path to an individually assigned vertex were assigned such a value. Thus, the two steps of the assignment procedure are feasible.

▀

**Theorem 4.24.** *If  $e$  is a consistent set, then the assignment  $a_e$  satisfies all the literals in  $e$ .*

*Proof.* Consider first the case of two variables  $x_i$  and  $x_j$  that are connected by a  $G_{\perp}^E(e)$ -edge. We have to show that  $a_e(x_i) = a_e(x_j)$ . Since  $x_i$  and  $x_j$  are  $G_{\perp}^E(e)$ -connected, they belong to the same  $G_{\perp}^E(e)$ -connected component. If they were both assigned a value by **R1**, then they were assigned the minimal value of an individually assigned vertex to which they are both  $G_{\perp}^E(e)$ -connected. If, on the other hand, they were both assigned a value by **R2**, then they were assigned the communal value assigned to the  $G_{\perp}^E$  component to which they both belong. Thus, in both cases they are assigned the same value.

Next, consider the case of two variables  $x_i$  and  $x_j$  that are connected by a  $G_{\neq}^E(e)$ -edge. To show that  $a_e(x_i) \neq a_e(x_j)$ , we distinguish three cases:

- If both  $x_i$  and  $x_j$  were assigned values by **R1**, they must have inherited their values from two distinct individually assigned vertices, because, otherwise, they are both connected by a  $G_{\perp}^E(e)$ -path to a common vertex, which together with the  $(x_i, x_j)$   $G_{\neq}^E(e)$ -edge closes a contradictory cycle, excluded by the assumption that  $e$  is consistent.

- If one of  $x_i, x_j$  was assigned a value by **R1** and the other acquired its value from **R2**, then since any communal value is distinct from any individually assigned value,  $a_e(x_i)$  must differ from  $a_e(x_j)$ .
- The remaining case is when both  $x_i$  and  $x_j$  were assigned values by **R2**. The fact that they were not assigned values in **R1** implies that their characteristic values are not individually allocated, but communally allocated. Assume falsely that  $a_e(x_i) = a_e(x_j)$ . This means that  $x_i$  and  $x_j$  were allocated their communal values in the same step, II.B, of the allocation algorithm, which implies that they had an equality path between them (moreover, this path was still part of the graph at the beginning of step II.B). Hence,  $x_i$  and  $x_j$  belong to a contradictory cycle, and the solid edge  $(x_i, x_j)$  was therefore still part of  $G_{\underline{=}}^E(e)$  at the beginning of step II.A. According to the loop condition of this step, at the end of this step there are no mixed vertices left, which rules out the possibility that  $(x_i, x_j)$  was still part of the graph at that stage. Thus, at least one of these vertices was individually assigned a value in step II.A.1, and, consequently, the component that it belongs to is assigned a value by **R1**, in contradiction to our assumption.  $\blacksquare$

**Theorem 4.25.** *The formula  $\varphi^E$  is satisfiable if and only if  $\varphi^E$  is satisfiable over  $D$ .*

*Proof.* By Theorems 4.23 and 4.24,  $D$  is adequate for  $E_{=} \cup E_{\neq}$ . Consequently,  $D$  is adequate for  $\Phi(\text{At}(\varphi^E))$ , and in particular  $D$  is adequate for  $\varphi^E$ . Thus, by the definition of adequacy,  $\varphi^E$  is satisfiable if and only if  $\varphi^E$  is satisfiable over  $D$ .  $\blacksquare$

#### 4.5.5 Summary

To summarize Sect. 4.5, the domain allocation method can be used as the first stage of a decision procedure for equality logic. In the second stage, the allocated domains can be enumerated by a standard BDD or by a SAT-based tool. Domain allocation has the advantage of not changing (in particular, not increasing) the original formula, unlike the algorithm that we studied in Sect. 4.4. Moreover, Algorithm 4.5.1 is highly effective in practice in allocating very small domains.

## 4.6 Ackermann's vs. Bryant's Reduction: Where Does It Matter?

We conclude this chapter by demonstrating how the two reductions lead to different equality graphs and hence change the result of applying any of the algorithms studied in this chapter that are based on this equality graph.

**Example 4.26.** Suppose that we want to check the satisfiability of the following (satisfiable) formula:

$$\varphi^{\text{UF}} := x_1 = x_2 \vee (F(x_1) \neq F(x_2) \wedge \text{FALSE}). \quad (4.31)$$

With Ackermann's reduction, we obtain:

$$\varphi^{\text{E}} := (x_1 = x_2 \implies f_1 = f_2) \wedge (x_1 = x_2 \vee (f_1 \neq f_2 \wedge \text{FALSE})). \quad (4.32)$$

With Bryant's reduction, we obtain:

$$\text{flat}^{\text{E}} := x_1 = x_2 \vee (F_1^* \neq F_2^* \wedge \text{FALSE}), \quad (4.33)$$

$$FC^{\text{E}} := \begin{array}{l} F_1^* = f_1 \\ F_2^* = \left( \begin{array}{l} \text{case } x_1 = x_2 : f_1 \\ \text{TRUE} : f_2 \end{array} \right), \end{array} \quad \wedge \quad (4.34)$$

and, as always,

$$\varphi^{\text{E}} := FC^{\text{E}} \wedge \text{flat}^{\text{E}}. \quad (4.35)$$

The equality graphs corresponding to the two reductions appear in Fig. 4.8. Clearly, the allocation for the right graph (due to Bryant's reduction) is smaller.



**Fig. 4.8.** The equality graph corresponding to Example 4.26 obtained with Ackermann's reduction (*left*) and with Bryant's reduction (*right*)

Indeed, an adequate range for the graph on the right is

$$D := \{x_1 \mapsto \{0\}, x_2 \mapsto \{0, 1\}, f_1 \mapsto \{2\}, f_2 \mapsto \{3\}\}. \quad (4.36)$$

These domains are adequate for (4.35), since we can choose the satisfying assignment

$$\{x_1 \mapsto 0, x_2 \mapsto 0, f_1 \mapsto 2, f_2 \mapsto 3\}. \quad (4.37)$$

On the other hand, this domain is not adequate for (4.32).

In order to satisfy (4.32), it must hold that  $x_1 = x_2$ , which implies that  $f_1 = f_2$  must hold as well. But the domains allocated in (4.36) do not allow an assignment in which  $f_1$  is equal to  $f_2$ , which means that the graph on the right of Fig. 4.8 is not adequate for (4.32).  $\blacksquare$

So what has happened here? Why does Ackermann's reduction require a larger range?

The reason is that when two function instances  $F(x_1)$  and  $F(x_2)$  have equal arguments, in Ackermann's reduction the two variables representing the functions, say  $f_1$  and  $f_2$ , are constrained to be equal. But if we force  $f_1$  and  $f_2$  to be different (by giving them a singleton domain composed of a unique constant), this forces  $FC^E$  to be FALSE, and, consequently  $\varphi^E$  to be FALSE. On the other hand, in Bryant's reduction, if the arguments  $x_1$  and  $x_2$  are equal, the terms  $F_1^*$  and  $F_2^*$  that represent the two functions are both assigned the value of  $f_1$ . Thus, even if  $f_2 \neq f_1$ , this does not necessarily make  $FC^E$  FALSE.

In the bibliographic notes of this chapter, we mention several publications that exploit this property of Bryant's reduction for reducing the allocated range and even constructing smaller equality graphs. It turns out that not all of the edges that are associated with the functional-consistency constraints are necessary, which, in turn, results in a smaller allocated range.

## 4.7 Problems

### 4.7.1 Conjunctions of Equalities and Uninterpreted Functions

**Problem 4.1 (deciding a conjunction of equalities with equivalence classes).** Consider Algorithm 4.7.1. Present details of an efficient implementation of this algorithm, including a data structure. What is the complexity of your implementation?

**Algorithm 4.7.1:** DECIDE-A-CONJUNCTION-OF-EQUALITIES-WITH-EQUIVALENCE-CLASSES

**Input:** A conjunction  $\varphi^E$  of equality predicates

**Output:** "Satisfiable" if  $\varphi^E$  is satisfiable, and "Unsatisfiable" otherwise

1. Define an equivalence class for each variable. For each equality  $x = y$  in  $\varphi^E$ , unite the equivalence classes of  $x$  and  $y$ .
2. For each disequality  $u \neq v$  in  $\varphi^E$ , if  $u$  is in the same equivalence class as  $v$ , return "Unsatisfiable".
3. Return "Satisfiable".

**Problem 4.2 (deciding a conjunction of equality predicates with a graph analysis).** Show a graph-based algorithm for deciding whether a given conjunction of equality predicates is satisfiable. What is the complexity of your algorithm?

**Problem 4.3 (a generalization of the CONGRUENCE-CLOSURE algorithm).** Generalize Algorithm 4.1.1 to the case in which the input formula includes uninterpreted functions with multiple arguments.

### 4.7.2 Reductions

**Problem 4.4 (a better way to eliminate constants?).** Is the following theorem correct?

**Theorem 4.27.** *An equality formula  $\varphi^E$  is satisfiable if and only if the formula  $\varphi^{E'}$  generated by Algorithm 4.7.2 (REMOVE-CONSTANTS-OPTIMIZED) is satisfiable.*

Prove the theorem or give a counterexample. You may use the result of Problem 3.2 in your proof.

**Algorithm 4.7.2: REMOVE-CONSTANTS-OPTIMIZED**

**Input:** An equality logic formula  $\varphi^E$


**Output:** An equality logic formula  $\varphi^{E'}$  such that  $\varphi^{E'}$  contains no constants and  $\varphi^{E'}$  is satisfiable if and only if  $\varphi^E$  is satisfiable

1.  $\varphi^{E'} := \varphi^E$ .
2. Replace each constant  $c$  in  $\varphi^{E'}$  with a new variable  $C_c$ .
3. For each pair of constants  $c_i, c_j$  with an equality path between them ( $c_i =^* c_j$ ) *not through any other constant*, add the constraint  $C_{c_i} \neq C_{c_j}$  to  $\varphi^{E'}$ . (Recall that the equality path is defined over  $G^E(\varphi^E)$ , where  $\varphi^E$  is given in NNF.)

**Problem 4.5 (correctness of the simplification step).** Prove the correctness of Algorithm 4.3.1. You may use the proof strategy suggested in Sect. 4.3.

**Problem 4.6 (reduced transitivity constraints).** (Based on [126, 169].) Consider the equality graph in Fig. 4.9. The sparse method generates  $\mathcal{B}_{trans}$  with three transitivity constraints (recall that it generates three constraints for each triangle in the graph, regardless of the polarity of the edges). Now consider the following claim: the single transitivity constraint  $\mathcal{B}_{rtc} = (e_{0,2} \wedge e_{1,2} \implies e_{0,1})$  is sufficient (the subscript *rtc* stands for “reduced transitivity constraints”).

To justify this claim, it is sufficient to show that for every assignment  $\alpha_{rtc}$  that satisfies  $e(\varphi^E) \wedge \mathcal{B}_{rtc}$ , there exists an assignment  $\alpha_{trans}$  that satisfies  $e(\varphi^E) \wedge \mathcal{B}_{trans}$ . Since this, in turn, implies that  $\varphi^E$  is satisfiable as well, we obtain the result that  $\varphi^E$  is satisfiable if and only if  $e(\varphi^E) \wedge \mathcal{B}_{rtc}$  is satisfiable.



	$\alpha_{rtc}$	$\alpha_{trans}$
$e_{0,1}$	TRUE	TRUE
$e_{1,2}$	TRUE	TRUE
$e_{0,2}$	FALSE	TRUE

**Fig. 4.9.** Taking polarity into account allows us to construct a less constrained formula. For this graph, the constraint  $\mathcal{B}_{rtc} = (e_{0,2} \wedge e_{1,2} \implies e_{0,1})$  is sufficient. An assignment  $\alpha_{rtc}$  that satisfies  $\mathcal{B}_{rtc}$  but breaks transitivity can always be “fixed” so that it *does* satisfy transitivity, while still satisfying the propositional skeleton  $e(\varphi^E)$ . The assignment  $\alpha_{trans}$  demonstrates such a “fixed” version of the satisfying assignment

We are able to construct such an assignment  $\alpha_{trans}$  because of the monotonicity of NNF (see Theorem 1.14, and recall that the polarity of the edges in the equality graph is defined according to their polarity in the NNF representation of  $\varphi^E$ ). There are only two satisfying assignments to  $\mathcal{B}_{rtc}$  that do not satisfy  $\mathcal{B}_{trans}$ . One of these assignments is shown in the  $\alpha_{rtc}$  column in the table to the right of the drawing. The second column shows a corresponding assignment  $\alpha_{trans}$ , which clearly satisfies  $\mathcal{B}_{trans}$ .

However, we still need to prove that every formula  $e(\varphi^E)$  that corresponds to the above graph is still satisfied by  $\alpha_{trans}$  if it is satisfied by  $\alpha_{rtc}$ . For example, for  $e(\varphi^E) = (\neg e_{0,1} \vee e_{1,2} \vee e_{0,2})$ , both  $\alpha_{rtc} \models e(\varphi^E) \wedge \mathcal{B}_{rtc}$  and  $\alpha_{trans} \models e(\varphi^E) \wedge \mathcal{B}_{trans}$ . Intuitively, this is guaranteed to be true because  $\alpha_{trans}$  is derived from  $\alpha_{rtc}$  by flipping an assignment of a positive (unnegated) predicate ( $e_{0,2}$ ) from FALSE to TRUE. We can equivalently flip an assignment to a negated predicate ( $e_{0,1}$  in this case) from TRUE to FALSE.

1. Generalize this example into a claim: given a (polar) equality graph, which transitivity constraints are necessary and sufficient?
2. Show an algorithm that computes the constraints that you suggest in the item above. What is the complexity of your algorithm? (*Hint*: there exists a polynomial algorithm, which is hard to find. An exponential algorithm will suffice as an answer to this question).

### 4.7.3 Complexity

**Problem 4.7 (complexity of deciding equality logic).** Prove that deciding equality logic is NP-complete.

Note that to show membership in NP, it is not enough to say that every solution can be checked in P-time, because the solution itself can be arbitrarily large, and hence even reading it is not necessarily a P-time operation.

#### 4.7.4 Domain Allocation

**Problem 4.8 (adequate domain for  $\Phi_{n,k}$ ).** Prove Theorem 4.20.

**Problem 4.9 (small-domain allocation).** Prove the following lemma.

**Lemma 4.28.** *If a domain  $D$  is adequate for  $\Phi(e)$  and  $e' \subseteq e$ , then  $D$  is adequate for  $\phi(e')$ .*

**Problem 4.10 (small-domain allocation: an adequate domain).** Prove the following theorem:

**Theorem 4.29.** *If all the subsets of  $E(\varphi^E)$  are consistent, then there exists an allocation  $R$  such that  $|R| = 1$ .*

**Problem 4.11 (formulation of the graph-theoretic problem).** Give a self-contained formal definition of the following decision problem: given an equality graph  $G$  and a domain allocation  $D$ , is  $D$  adequate for  $G$ ?

**Problem 4.12 (small-domain allocation: an improvement to the allocation heuristic).** Step II.A.1 of Algorithm 4.5.1 calls for allocation of *distinct* characteristic values to the mixed vertices. The following example proves that this is not always necessary.

Consider the subgraph  $\{u_1, f_1, f_2, u_2\}$  of the graph in Fig. 4.2. Application of the basic algorithm to this subgraph may yield the following allocation, where the characteristic values assigned are underlined:  $R_1 : u_1 \mapsto \{0, \underline{2}\}, f_1 \mapsto \{\underline{0}\}, f_2 \mapsto \{0, \underline{1}\}, u_2 \mapsto \{0, 1, \underline{3}\}$ . This allocation leads to a state space complexity of 12. By relaxing the requirement that all individually assigned characteristic values should be distinct, we can obtain the allocation  $R_2 : u_1 \mapsto \{0, \underline{2}\}, f_1 \mapsto \{\underline{0}\}, f_2 \mapsto \{\underline{0}\}, u_2 \mapsto \{0, \underline{1}\}$  with a state-space complexity of 4. This reduces the size of the state space of the entire graph from 48 to 16.

It is not difficult to see that  $R_2$  is adequate for the subgraph considered.

What are the conditions under which it is possible to assign equal values to mixed variables? Change the basic algorithm so that it includes this optimization.

## 4.8 Bibliographic Notes

The treatment of equalities and uninterpreted functions can be divided into several eras. In the first era, before the emergence of the first effective theorem provers in the 1970's, this logic was considered only from the point of view of mathematical logic, most notably by Ackermann [1]. In the same book, he also offered what we have called Ackermann's reduction in this book. Equalities were typically handled with rewriting rules, for example substituting  $x$  with  $y$  given that  $x = y$ .

The second era started in the mid 1970's with the work of Downey, Sethi, and Tarjan [69], who showed that the decision problem was a variation on the common-subexpression problem; the work of Nelson and Oppen [136], who applied the union-find algorithm to compute the congruence closure and implemented it in the Stanford Pascal Verifier; and then the work of Shostak, who suggested in [178] the congruence closure method that was briefly presented in Sect. 4.1. All of this work was based on computing the congruence closure, and indicated a shift from the previous era, as it offered complete and relatively efficient methods for deciding equalities and uninterpreted functions. In its original presentation, Shostak's method relied on syntactic case-splitting (see Sect. 1.3), which is the source of the inefficiency of that algorithm. In Shostak's words, "it was found that most examples four or five lines long could be handled in just a few seconds". Even factoring in the fact that this was done on a 1978 computer (a DEC-10 computer), this statement still shows how much progress has been made since then, as nowadays many formulas with tens of thousands of variables are solved in a few seconds. Several variants on Shostak's method exist, and have been compared and described in a single theoretical framework called **abstract congruence closure** in [8]. Shostak's method and its variants are still used in theorem provers, although several improvements have been suggested to combat the practical complexity of case-splitting, namely *lazy case-splitting*, in which the formula is split only when it is necessary for the proof, and other similar techniques.

The third era of deciding this theory avoided syntactic case-splitting altogether and instead promoted the use of *semantic* case-splitting, that is, splitting the domain instead of splitting the formula. All of the methods of this type are based on an underlying decision procedure for Boolean formulas, such as a SAT engine or the use of BDDs. We failed to find an original reference for the fact that the range  $\{1, \dots, n\}$  is adequate for formulas with  $n$  variables. This is usually referred to as a "folk theorem" in the literature. The work by Hojati, Kuehlmann, German, and Brayton in [95] and Hojati, Isles, Kirkpatrick, and Brayton in [94] was the first, as far as we know, where anyone tried to decide equalities with finite instantiation, while trying to derive a value  $k$ ,  $k \leq n$  that was adequate as well, by analyzing the equality graph. The method presented in Sect. 4.5 was the first to consider a different range for each variable and, hence, is much more effective. It is based on work by Pnueli, Rodeh, Siegel, and Strichman in [154, 155]. These papers suggest that Ackermann's reduction should be used, which results in large formulas, and, consequently, large equality graphs and correspondingly large domains (but much smaller than the range  $\{1, \dots, n\}$ ). Bryant, German and Velev suggested in [38] what we refer to as Bryant's reduction in Sect. 3.3.2. This technique enabled them to exploit what they called the *positive equality* structure in formulas for assigning unique constants to some of the variables and a full range to the others. Using the terminology of this chapter, these variables are adjacent only to solid edges in the equality graph corresponding to the original formula (a graph built *without* referring to the functional-consistency



constraints, and hence the problem of a large graph due to Ackermann's constraints disappears). A more robust version of this technique, in which a larger set of variables can be replaced with constants, was later developed by Lahiri, Bryant, Goel, and Talupur [112].

In [167, 168], Rodeh and Shtrichman presented a generalization of positive equality that enjoys benefits from both worlds: on the one hand, it does not add all the edges that are associated with the functional-consistency constraints (it adds only a small subset of them based on an analysis of the formula), but on the other hand it assigns small ranges to *all* variables as in [155] and, in particular, a single value to all the terms that would be assigned a single value by the technique of [38]. This method decreases the size of the equality graph in the presence of uninterpreted functions, and consequently the allocated ranges (for example, it allocates a domain with a state space of size 2 for the running example in Sect. 4.5.3). Rodeh showed in his thesis [167] (also see [153]) an extension of range allocation to *dynamic* range allocation. This means that each variable is assigned not one of several constants, as prescribed by the allocated domain, but rather one of the variables that represent an immediate neighbor in  $G_{=}^E$ , or a unique constant if it has one or more neighbors in  $G_{\neq}^E$ . The size of the state space is thus proportional to  $\log n$ , where  $n$  is the number of neighbors.

Goel, Sajid, Zhou, Aziz, and Singhal were the first to encode each equality with a new Boolean variable [87]. They built a BDD corresponding to the encoded formula, and then looked for transitivity-preserving paths in the BDD. Bryant and Velev suggested in [39] that the same encoding should be used but added explicit transitivity constraints instead. They considered several translation methods, only the best of which (the sparse method) was presented in this chapter. One of the other alternatives is to add such a constraint for every three variables (regardless of the equality graph). A somewhat similar approach was considered by Zantema and Groote [206]. The sparse method was later superseded by the method of Meir and Strichman [126] and later by that of Rozanov and Strichman [169], where the polar equality graph is considered rather than the nonpolar one, which leads to a smaller number of transitivity constraints. This direction is mentioned in Problem 4.6.

All the methods that we discussed in this chapter, other than congruence closure, belong to the third era. A fourth era, based on an interplay between a SAT solver and a decision procedure for a conjunction of terms (such as congruence closure in the case of EUF formulas), has emerged in the last few years, and is described in detail in Chap. 11. The idea is also explained briefly at the end of Sect. 4.1.

## 4.9 Glossary

The following symbols were used in this chapter:

<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$E_-, E_{\neq}$	Sets of equality and inequality predicates, and also the edges in the equality graph	83
$At(\varphi^E)$	The set of atoms in the formula $\varphi^E$	83
$G^E$	Equality graph	84
$x =^* y$	There exists an equality path between $x$ and $y$ in the equality graph	84
$x \neq^* y$	There exists a disequality path between $x$ and $y$ in the equality graph	84
$e(\varphi^E)$	The propositional skeleton of $\varphi^E$	89
$\mathcal{B}_{trans}$	The transitivity constraints due to the reduction from $\varphi^E$ to $\mathcal{B}_{sat}$ by the sparse method	89
$G_{NP}^E$	Nonpolar equality graph	89
$var(\varphi^E)$	The set of variables in $\varphi^E$	92
$D$	A domain allocation function. See (4.23)	92
$ D $	The state space spanned by a domain	92
$\Phi_n$	The (infinite) set of equality logic formulas with $n$ variables	93
$\Phi_{n,k}$	The (infinite) set of equality logic formulas with $n$ variables and $k$ constants	94
$\phi(e)$	The (infinite) set of equality formulas with a set of literals equal to $e$	95
$E(\varphi^E)$	The set of literals in $\varphi^E$	95
$G_-, G_{\neq}^E$	The projections of the equality graph on the $E_-$ and $E_{\neq}$ edges, respectively	96
$char(v)$	The characteristic value of a node $v$ in the equality graph	96
<i>continued on next page</i>		

*continued from previous page*

<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$MV$	The set of mixed vertices that are chosen in step II.A.1 of Algorithm 4.5.1	96
$a_e(x)$	An assignment to a variable $x$ from its allocated domain $D(x)$	99

---

## Linear Arithmetic

### 5.1 Introduction

This chapter introduces decision procedures for conjunctions of linear constraints. An extension of these decision procedures for solving a general linear arithmetic formula, i.e., with an arbitrary Boolean structure, is given in Chap. 11.

**Definition 5.1 (linear arithmetic).** *The syntax of a formula in linear arithmetic is defined by the following rules:*

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid (\text{formula}) \mid \text{atom} \\ \text{atom} &: \text{sum} \text{ op } \text{sum} \\ \text{op} &: = \mid \leq \mid < \\ \text{sum} &: \text{term} \mid \text{sum} + \text{term} \\ \text{term} &: \text{identifier} \mid \text{constant} \mid \text{constant identifier} \end{aligned}$$

The binary minus operator  $a - b$  can be read as “syntactic sugar” for  $a + -1b$ . The operators  $\geq$  and  $>$  can be replaced by  $\leq$  and  $<$  if the coefficients are negated. We consider the rational numbers and the integers as domains. For the former domain the problem is polynomial, and for the latter the problem is NP-complete.

As an example, the following is a formula in linear arithmetic:

$$3x_1 + 2x_2 \leq 5x_3 \quad \wedge \quad 2x_1 - 2x_2 = 0. \quad (5.1)$$

Note that equality logic, as discussed in Chap. 4, is a fragment of linear arithmetic.

Many problems arising in the code optimization performed by compilers are expressible with linear arithmetic over the integers. As an example, consider the following C code fragment:

```

for(i=1; i<=10; i++)
  a[j+i]=a[j];

```

This fragment is intended to replicate the value of  $a[j]$  into the locations  $a[j+1]$  to  $a[j+10]$ . In a DLX-like assembly language,<sup>1</sup> a compiler might generate the code for the body of the loop as follows. Suppose variable  $i$  is stored in register R1, and variable  $j$  is stored in register R2:

```

R4 ← mem[a+R2]      /* set R4 to a[j] */
R5 ← R2+R1          /* set R5 to j+i */
mem[a+R5] ← R4      /* set a[j+i] to a[j] */
R1 ← R1+1           /* i++ */

```

Code that requires memory access is typically very slow compared with code that operates only on the internal registers of the CPU. Thus, it is highly desirable to avoid load and store instructions. A potential optimization for the code above is to move the load instruction for  $a[j]$ , i.e., the first statement above, out of the loop body. After this transformation, the load instruction is executed only once at the beginning of the loop, instead of 10 times. However, the correctness of this transformation relies on the fact that the value of  $a[j]$  does not change within the loop body. We can check this condition by comparing the index of  $a[j+i]$  with the index of  $a[j]$  together with the constraint that  $i$  is between 1 and 10:

$$i \geq 1 \wedge i \leq 10 \wedge j + i = j. \quad (5.2)$$

This formula has no satisfying assignment, and thus, the memory accesses cannot overlap. The compiler can safely perform the read access to  $a[j]$  only once.

### 5.1.1 Solvers for Linear Arithmetic

The **simplex** method is one of the oldest algorithms for numerical optimization. It is used to find an optimal value for an objective function given a conjunction of linear constraints over real variables. The objective function and the constraints together are called a **linear program** (LP). However, since we are interested in the decision problem rather than the optimization problem, we cover in this chapter a variant of the simplex method called **general simplex** that takes as input a conjunction of linear constraints over the reals *without* an objective function, and decides whether this set is satisfiable.

**Integer linear programming**, or ILP, is the same problem for constraints over integers. Section 5.3 covers BRANCH AND BOUND, an algorithm for deciding such problems.

<sup>1</sup> The DLX architecture is a RISC-like computer architecture, which is similar to the MIPS architecture [149].

These two algorithms can solve conjunctions of a large number of constraints efficiently. We shall also describe two other methods that are considered less efficient, but can still be competitive for solving small problems. We describe them because they are still used in practice, they are relatively easy to implement in their basic form, and they will be mentioned again later in Chap. 11, owing to the fact that they are based on variable elimination. The first of these methods is called **Fourier–Motzkin** variable elimination, and decides the satisfiability of a conjunction of linear constraints over the reals. The second method is called **Omega test**, and decides the satisfiability of a conjunction of linear constraints over the integers.

## 5.2 The Simplex Algorithm

### 5.2.1 Decision Problems and Linear Programs

The simplex algorithm, originally developed by Danzig in 1947, decides satisfiability of a conjunction of weak linear inequalities. The set of constraints is normally accompanied by a linear *objective function* in terms of the variables of the formula. If the set of constraints is satisfiable, the simplex algorithm provides a satisfying assignment that maximizes the value of the objective function. Simplex is worst-case exponential. Although there are polynomial-time algorithms for solving this problem (the first known polynomial-time algorithm, introduced by Khachiyan in 1979, is called the **ellipsoid method**), simplex is still considered a very efficient method in practice and the most widely used, apparently because the need for an exponential number of steps is rare in real problems.

As we are concerned with the decision problem rather than the optimization problem, we are going to cover a variant of the simplex algorithm called **general simplex** that does not require an objective function. The general simplex algorithm accepts two types of constraints as input:

1. Equalities of the form

$$a_1x_1 + \dots + a_nx_n = 0. \quad (5.3)$$

2. Lower and upper bounds on the variables:<sup>2</sup>

$$l_i \leq x_i \leq u_i, \quad (5.4)$$

where  $l_i$  and  $u_i$  are constants representing the lower and upper bounds on  $x_i$ , respectively. The bounds are optional as the algorithm supports unbounded variables.

 $l_i$ 
 $u_i$ 

<sup>2</sup> This is in contrast to the classical simplex algorithm, in which all variables are constrained to be nonnegative.

This representation of the input formula is called the **general form**. This statement of the problem does not restrict the modeling power of weak linear constraints, as we can transform an arbitrary weak linear constraint  $L \bowtie R$  with  $\bowtie \in \{=, \leq, \geq\}$  into the form above as follows. Let  $m$  be the number of constraints. For the  $i$ -th constraint,  $1 \leq i \leq m$ :

1. Move all addends in  $R$  to the left-hand side to obtain  $L' \bowtie b$ , where  $b$  is a constant.
2. Introduce a new variable  $s_i$ . Add the constraints

$$L' - s_i = 0 \quad \text{and} \quad s_i \bowtie b. \quad (5.5)$$

If  $\bowtie$  is the equality operator, rewrite  $s_i = b$  to  $s_i \geq b$  and  $s_i \leq b$ .

The original and the transformed conjunctions of constraints are obviously equisatisfiable.

**Example 5.2.** Consider the following conjunction of constraints:

$$\begin{aligned} x + y &\geq 2 \wedge \\ 2x - y &\geq 0 \wedge \\ -x + 2y &\geq 1 \quad . \end{aligned} \quad (5.6)$$

The problem is rewritten into the general form as follows:

$$\begin{aligned} x + y - s_1 &= 0 \wedge \\ 2x - y - s_2 &= 0 \wedge \\ -x + 2y - s_3 &= 0 \wedge \\ s_1 &\geq 2 \wedge \\ s_2 &\geq 0 \wedge \\ s_3 &\geq 1 \quad . \end{aligned} \quad (5.7)$$


The new variables  $s_1, \dots, s_m$  are called the **additional variables**. The variables  $x_1, \dots, x_n$  in the original constraints are called **problem variables**. Thus, we have  $n$  problem variables and  $m$  additional variables. As an optimization of the procedure above, an additional variable is only introduced if  $L'$  is not already a problem variable or has been assigned an additional variable previously.

### 5.2.2 Basics of the Simplex Algorithm

It is common and convenient to view linear constraint satisfaction problems as geometrical problems. In geometrical terms, each variable corresponds to a dimension, and each constraint defines a convex subspace: in particular, inequalities define *half-spaces* and equalities define hyperplanes.<sup>3</sup> The (closed)

<sup>3</sup> A hyperplane in a  $d$ -dimensional space is a subspace with  $d - 1$  dimensions. For example, in two dimensions, a hyperplane is a straight line, and in one dimension it is a point.

subspace of satisfying assignments is defined by an intersection of half spaces and hyperplanes, and forms a convex polytope. This is implied by the fact that an intersection between convex subspaces is convex as well. A geometrical representation of the original problem in Example 5.2 appears in Fig. 5.1.



**Fig. 5.1.** A graphical representation of the problem in Example 5.2, projected on  $x$  and  $y$ . The shaded region corresponds to the set of satisfying assignments. The marked points (A), (B), and (C) illustrate the progress that the simplex algorithm makes, as will be explained in the rest of this section

It is common to represent the coefficients in the input problem using an  $m$ -by- $(n + m)$  matrix  $A$ . The variables  $x_1, \dots, x_n, s_1, \dots, s_m$  are written as a vector  $\mathbf{x}$ . Following this notation, our problem is equivalent to the existence of a vector  $\mathbf{x}$  such that

$$A\mathbf{x} = \mathbf{0} \quad \text{and} \quad \bigwedge_{i=1}^m l_i \leq s_i \leq u_i, \quad (5.8)$$

where  $l_i \in \{-\infty\} \cup \mathbb{Q}$  is the lower bound of  $x_i$  and  $u_i \in \{+\infty\} \cup \mathbb{Q}$  is the upper bound of  $x_i$ . The infinity values are for the case that a bound is not set.

**Example 5.3.** We continue Example 5.2. Using the variable ordering  $x, y, s_1, s_2, s_3$ , a matrix representation for the equality constraints in (5.7) is

$$\begin{pmatrix} 1 & 1 & -1 & 0 & 0 \\ 2 & -1 & 0 & -1 & 0 \\ -1 & 2 & 0 & 0 & -1 \end{pmatrix}. \quad (5.9)$$

▀

Note that a large portion of the matrix in Example 5.3 is very regular: the columns that are added for the additional variables  $s_1, \dots, s_m$  correspond to



an  $m$ -by- $m$  diagonal matrix, where the diagonal coefficients are  $-1$ . This is a direct consequence of using the general form.

While the matrix  $A$  changes as the algorithm progresses, the number of columns of this kind is never reduced. The set of  $m$  variables corresponding to these columns are called the **basic variables** and denoted by  $\mathcal{B}$ . They are also called the *dependent* variables, as their values are determined by those of the nonbasic variables. The nonbasic variables are denoted by  $\mathcal{N}$ . It is convenient to store and manipulate a representation of  $A$  called the **tableau**, which is simply  $A$  without the diagonal submatrix. The tableau is thus an  $m$ -by- $n$  matrix, where the columns correspond to the nonbasic variables, and each row is associated with a basic variable – the same basic variable that has a “ $-1$ ” entry at that row in the diagonal sub-matrix in  $A$ . Thus, the information originally stored in the diagonal matrix is now represented by the variables labeling the rows.

$\mathcal{B}, \mathcal{N}$

**Example 5.4.** We continue our running example. The tableau and the bounds for Example 5.2 are:

	$x$	$y$	
$s_1$	1	1	$2 \leq s_1$
$s_2$	2	$-1$	$0 \leq s_2$
$s_3$	$-1$	2	$1 \leq s_3$

▀

The tableau is simply a different representation of  $A$ , since  $Ax = 0$  can be rewritten into

$$\bigwedge_{x_i \in \mathcal{B}} \left( x_i = \sum_{x_j \in \mathcal{N}} a_{ij} x_j \right). \quad (5.10)$$

When written in the form of a matrix, the sums on the right-hand side of (5.10) correspond exactly to the tableau.

### 5.2.3 Simplex with Upper and Lower Bounds

The general simplex algorithm maintains, in addition to the tableau, an assignment  $\alpha : \mathcal{B} \cup \mathcal{N} \rightarrow \mathbb{Q}$ . The algorithm initializes its data structures as follows:

$\alpha$

- The set of basic variables  $\mathcal{B}$  is the set of additional variables.
- The set of nonbasic variables  $\mathcal{N}$  is the set of problem variables.
- For any  $x_i$  with  $i \in \{1, \dots, n + m\}$ ,  $\alpha(x_i) = 0$ .

If the initial assignment of zero to all variables (i.e., the origin) satisfies all upper and lower bounds of the basic variables, then the formula can be declared satisfiable (recall that initially the nonbasic variables do not have

**Algorithm 5.2.1:** GENERAL-SIMPLEX**Input:** A linear system of constraints  $S$ **Output:** “Satisfiable” if the system is satisfiable, “Unsatisfiable” otherwise

1. Transform the system into the general form

$$A\mathbf{x} = 0 \quad \text{and} \quad \bigwedge_{i=1}^m l_i \leq s_i \leq u_i .$$

2. Set  $\mathcal{B}$  to be the set of additional variables  $s_1, \dots, s_m$ .
3. Construct the tableau for  $A$ .
4. Determine a fixed order on the variables.
5. If there is no basic variable that violates its bounds, return “Satisfiable”. Otherwise, let  $x_i$  be the first basic variable in the order that violates its bounds.
6. Search for the first suitable nonbasic variable  $x_j$  in the order for pivoting with  $x_i$ . If there is no such variable, return “Unsatisfiable”.
7. Perform the pivot operation on  $x_i$  and  $x_j$ .
8. Go to step 5.

explicit bounds). Otherwise, the algorithm begins a process of changing this assignment.

Algorithm 5.2.1 summarizes the steps of the general simplex procedure. The algorithm maintains two invariants:

- **In-1.**  $A\mathbf{x} = 0$
- **In-2.** The values of the nonbasic variables are within their bounds:

$$\forall x_j \in \mathcal{N}. l_j \leq \alpha(x_j) \leq u_j . \quad (5.11)$$

Clearly, these invariants hold initially since all the variables in  $\mathbf{x}$  are set to 0, and the nonbasic variables have no bounds.

The main loop of the algorithm checks if there exists a basic variable that violates its bounds. If there is no such variable, then both the basic and nonbasic variables satisfy their bounds. Owing to invariant **In-1**, this means that the current assignment  $\alpha$  satisfies (5.8), and the algorithm returns “Satisfiable”.

Otherwise, let  $x_i$  be a basic variable that violates its bounds, and assume, without loss of generality, that  $\alpha(x_i) > u_i$ , i.e., the upper bound of  $x_i$  is violated. How do we change the assignment to  $x_i$  so it satisfies its bounds? We need to find a way to reduce the value of  $x_i$ . Recall how this value is specified:

$$x_i = \sum_{x_j \in \mathcal{N}} a_{ij} x_j. \quad (5.12)$$

The value of  $x_i$  can be reduced by decreasing the value of a nonbasic variable  $x_j$  such that  $a_{ij} > 0$  and its current assignment is higher than its lower bound  $l_j$ , or by increasing the value of a variable  $x_j$  such that  $a_{ij} < 0$  and its current assignment is lower than its upper bound  $u_j$ . A variable  $x_j$  fulfilling one of these conditions is said to be *suitable*. If there are no suitable variables, then the problem is unsatisfiable and the algorithm terminates.

$\theta$  Let  $\theta$  denote by how much we have to increase (or decrease)  $\alpha(x_j)$  in order to meet  $x_i$ 's upper bound:

$$\theta \doteq \frac{u_i - \alpha(x_i)}{a_{ij}}. \quad (5.13)$$

Increasing (or decreasing)  $x_j$  by  $\theta$  puts  $x_i$  within its bounds. On the other hand  $x_j$  does not necessarily satisfy its bounds anymore, and hence may violate the invariant **In-2**. We therefore swap  $x_i$  and  $x_j$  in the tableau, i.e., make  $x_i$  nonbasic and  $x_j$  basic. This requires a transformation of the tableau, which is called the **pivot operation**. The pivot operation is repeated until either a satisfying assignment is found, or the system is determined to be unsatisfiable.

### The Pivot Operation

Suppose we want to swap  $x_i$  with  $x_j$ . We will need the following definition:

**Definition 5.5 (pivot element, column and row).** *Given two variables  $x_i$  and  $x_j$ , the coefficient  $a_{ij}$  is called the pivot element. The column of  $x_j$  is called the pivot column. The row  $i$  is called the pivot row.*

A precondition for swapping two variables  $x_i$  and  $x_j$  is that their pivot element is nonzero, i.e.,  $a_{ij} \neq 0$ . The pivot operation (or **pivoting**) is performed as follows:

1. Solve row  $i$  for  $x_j$ .
2. For all rows  $l \neq i$ , eliminate  $x_j$  by using the equality for  $x_j$  obtained from row  $i$ .

The reader may observe that the pivot operation is also the basic operation in the well-known **Gaussian variable elimination** procedure.

**Example 5.6.** We continue our running example. As described above, we initialize  $\alpha(x_i) = 0$ . This corresponds to point (A) in Fig. 5.1. Recall the tableau and the bounds:

	$x$	$y$	
$s_1$	1	1	$2 \leq s_1$
$s_2$	2	-1	$0 \leq s_2$
$s_3$	-1	2	$1 \leq s_3$

The lower bound of  $s_1$  is 2, which is violated. The nonbasic variable that is the lowest in the ordering is  $x$ . The variable  $x$  has a positive coefficient, but no upper bound, and is therefore suitable for the pivot operation. We need to increase  $s_1$  by 2 in order to meet the lower bound, which means that  $x$  has to increase by 2 as well ( $\theta = 2$ ). The first step of the pivot operation is to solve the row of  $s_1$  for  $x$ :

$$s_1 = x + y \iff x = s_1 - y. \quad (5.14)$$

This equality is now used to replace  $x$  in the other two rows:

$$s_2 = 2(s_1 - y) - y \iff s_2 = 2s_1 - 3y \quad (5.15)$$

$$s_3 = -(s_1 - y) + 2y \iff s_3 = -s_1 + 3y \quad (5.16)$$

Written as a tableau, the result of the pivot operation is:

	$s_1$	$y$	$\alpha(x) = 2$
$x$	1	-1	$\alpha(y) = 0$
$s_2$	2	-3	$\alpha(s_1) = 2$
$s_3$	-1	3	$\alpha(s_2) = 4$
			$\alpha(s_3) = -2$

This new state corresponds to point (B) in Fig. 5.1.

The lower bound of  $s_3$  is violated; this is the next basic variable that is selected. The only suitable variable for pivoting is  $y$ . We need to add 3 to  $s_3$  in order to meet the lower bound. This translates into

$$\theta = \frac{1 - (-2)}{3} = 1. \quad (5.17)$$

After performing the pivot operation with  $s_3$  and  $y$ , the final tableau is:

	$s_1$	$s_3$	$\alpha(x) = 1$
$x$	2/3	-1/3	$\alpha(y) = 1$
$s_2$	1	-1	$\alpha(s_1) = 2$
$y$	1/3	1/3	$\alpha(s_2) = 1$
			$\alpha(s_3) = 1$

This assignment  $\alpha$  satisfies the bounds, and thus  $\{x \mapsto 1, y \mapsto 1\}$  is a satisfying assignment. It corresponds to point (C) in Fig. 5.1. ▀

Selecting the pivot element according to a fixed ordering for the basic and nonbasic variable ensures that no set of basic variables is ever repeated, and hence guarantees termination (no cycling can occur). For a detailed proof see [71]. This way of selecting a pivot element is called **Bland's rule**.

### 5.2.4 Incremental Problems

Decision problems are often constructed in an **incremental** manner, that is, the formula is strengthened with additional conjuncts. This can make a once satisfiable formula unsatisfiable. One scenario in which an incremental decision procedure is useful is the DPLL(T) framework, which we study in Chap. 11.

The general simplex algorithm is well-suited for incremental problems. First, notice that any constraint can be disabled by removing its corresponding upper and lower bounds. The equality in the tableau is afterwards redundant, but will not render a satisfiable formula unsatisfiable. Second, the pivot operation performed on the tableau is an equivalence transformation, i.e., it preserves the set of solutions. We can therefore start the procedure with the tableau we have obtained from the previous set of bounds.

The addition of upper and lower bounds is implemented as follows:

- If a bound for a nonbasic variable was added, update the values of the nonbasic variables according to the tableau to restore **In-2**.
- Call Algorithm 5.2.1 to determine if the new problem is satisfiable. Start with step 5.

Furthermore, it is often desirable to *remove* constraints after they have been added. This is also relevant in the context of DPLL(T) because this algorithm activates and deactivates constraints. Normally constraints (or rather bounds) are removed when the current set of constraints is unsatisfiable. After removing a constraint the assignment has to be restored to a point at which it satisfied the two invariants of the general simplex algorithm. This can be done by simply restoring the assignment  $\alpha$  to the last known satisfying assignment. There is no need to modify the tableau.

## 5.3 The Branch and Bound Method

BRANCH AND BOUND is a widely used method for solving integer linear programs. As in the case of the simplex algorithm, BRANCH AND BOUND was developed for solving the optimization problem, but the description here focuses on an adaptation of this algorithm to the decision problem.

The integer linear systems considered here have the same form as described in Sect. 5.2, with the additional requirement that the value of any variable in a satisfying assignment must be drawn from the set of integers. Observe that it is easy to support strict inequalities simply by adding 1 to or subtracting 1 from the constant on the right-hand side.

**Definition 5.7 (relaxed problem).** *Given an integer linear system  $S$ , its relaxation is  $S$  without the integrality requirement (i.e., the variables are not required to be integer).*

We denote the relaxed problem of  $S$  by  $\text{relaxed}(S)$ . Assume the existence of a procedure  $LP_{feasible}$ , which receives a linear system  $S$  as input, and returns “Unsatisfiable” if  $S$  is unsatisfiable and a satisfying assignment otherwise.  $LP_{feasible}$  can be implemented with, for example, a variation of GENERAL-SIMPLEX (Algorithm 5.2.1) that outputs a satisfying assignment if  $S$  is satisfiable. Using these notions, Algorithm 5.3.1 decides an integer linear system of constraints (recall that only conjunctions of constraints are considered here).

**Algorithm 5.3.1:** FEASIBILITY-BRANCH-AND-BOUND

**Input:** An integer linear system  $S$

**Output:** “Satisfiable” if  $S$  is satisfiable, “Unsatisfiable” otherwise

```

1. procedure SEARCH-INTEGRAL-SOLUTION( $S$ )
2.    $res = LP_{feasible}(\text{relaxed}(S))$ ;
3.   if  $res = \text{“Unsatisfiable”}$  then return ;           ▷ prune branch
4.   else
5.     if  $res$  is integral then                       ▷ integer solution found
6.       abort(“Satisfiable”);
7.     else
8.       Select a variable  $v$  that is assigned a nonintegral value  $r$ ;
9.       SEARCH-INTEGRAL-SOLUTION ( $S \cup (v \leq \lfloor r \rfloor)$ );
10.      SEARCH-INTEGRAL-SOLUTION ( $S \cup (v \geq \lceil r \rceil)$ );
11.      ▷ no integer solution in this branch

11. procedure FEASIBILITY-BRANCH-AND-BOUND( $S$ )
12.   SEARCH-INTEGRAL-SOLUTION( $S$ );
13.   return (“Unsatisfiable”);

```

The idea of the algorithm is simple: it solves the relaxed problem with  $LP_{feasible}$ ; if the relaxed problem is unsatisfiable, it backtracks because there is also no integer solution in this branch. If, on the other hand, the relaxed problem is satisfiable and the solution returned by  $LP_{feasible}$  happens to be integral, it terminates – a satisfying integral solution has been found. Otherwise, the problem is split into two subproblems, which are then processed with a recursive call. The nature of this split is best illustrated by an example.

**Example 5.8.** Let  $x_1, \dots, x_4$  be the variables of  $S$ . Assume that  $LP_{feasible}$  returns the solution

$$(1, 0.7, 2.5, 3) \tag{5.18}$$

in line 2. In line 7, SEARCH-INTEGRAL-SOLUTION chooses between  $x_2$  and  $x_3$ , which are the variables that were assigned a nonintegral value. Suppose that

$x_2$  is chosen. In line 8,  $S$  (the linear system solved at the current recursion level) is then augmented with the constraint

$$x_2 \leq 0 \tag{5.19}$$

and sent for solving at a deeper recursion level. If no solution is found in this branch,  $S$  is augmented instead with

$$x_2 \geq 1 \tag{5.20}$$

and, once again, is sent to a deeper recursion level. If both these calls return, this implies that  $S$  has no satisfying solution, and hence the procedure returns (backtracks). Note that returning from the initial recursion level causes the calling function FEASIBILITY-BRANCH-AND-BOUND to return “Unsatisfiable”. ▀

Algorithm 5.3.1 is not complete: there are cases for which it will branch forever. As noted in [71], the system  $1 \leq 3x - 3y \leq 2$ , for example, has no integer solutions but unbounded real solutions, and causes the basic Branch and Bound algorithm to loop forever. In order to make the algorithm complete, it is necessary to rely on the small-model property that such formulas have (we used this property earlier in Sect. 4.5). Recall that this means that if there is a satisfying solution, then there is also such a solution within a finite bound, which, for this theory, is also computable. This means that once we have computed this bound on the domain of each variable, we can stop searching for a solution once we have passed it. A detailed study of this bound in the context of optimization problems can be found in [139]. The same bounds are applicable to the feasibility problem as well. Briefly, it was shown in [139] that given an integer linear system  $S$  with an  $M \times N$  coefficient matrix  $A$ , then if there is a solution to  $S$ , then one of the extreme points of the convex hull of  $S$  is also a solution, and any such solution  $x^0$  is bounded as follows:

$$x_j^0 \leq ((M + N) \cdot N \cdot \theta)^N \quad \text{for } j = 1, \dots, N, \tag{5.21}$$

where  $\theta$  is the maximal element in the coefficient matrix  $A$  or in the vector  $\mathbf{b}$ . Thus, (5.21) gives us a bound on each of the  $N$  variables, which, by adding it as an explicit constraint, forces termination.


Finally, let us mention that BRANCH AND BOUND can be extended in a straightforward way to handle the case in which some of the variables are integers while the others are real. In the context of optimization problems, this problem is known by the name **mixed integer programming**.

### 5.3.1 Cutting-Planes

*Cutting-planes* are constraints that are added to a linear system that remove only noninteger solutions; that is, all satisfying integer solutions, if they exist,

**Aside: BRANCH AND BOUND for Integer Linear Programs**

When BRANCH AND BOUND is used for solving an optimization problem, it becomes somewhat more complicated. In particular, there are various pruning rules based on the value of the current objective function (a branch is pruned if it is identified that it cannot contain a solution better than what is already at hand from another branch). There are also various heuristics for choosing the variable on which to split and the first branch to be explored.



**Fig. 5.2.** The dots represent integer solutions. The thin dotted line represents a cutting-plane – a constraint that does not remove any integral solution

remain satisfying, as demonstrated in Fig. 5.2. These new constraints improve the tightness of the relaxation in the process of solving integer linear systems.

Here, we describe a family of cutting planes called **Gomory cuts**. We first illustrate this technique with an example, and then generalize it.

Suppose that our problem includes the integer variables  $x_1, \dots, x_3$ , and the lower bounds  $1 \leq x_1$  and  $0.5 \leq x_2$ . Further, suppose that the final tableau of the general simplex algorithm includes the constraint

$$x_3 = 0.5x_1 + 2.5x_2, \quad (5.22)$$

and that the solution  $\alpha$  is  $\{x_3 \mapsto 1.75, x_1 \mapsto 1, x_2 \mapsto 0.5\}$ , which, of course, satisfies (5.22). Subtracting these values from (5.22) gives us

$$x_3 - 1.75 = 0.5(x_1 - 1) + 2.5(x_2 - 0.5). \quad (5.23)$$

We now wish to rewrite this equation so the left-hand side is an integer:

$$x_3 - 1 = 0.75 + 0.5(x_1 - 1) + 2.5(x_2 - 0.5). \quad (5.24)$$



The two right-most terms must be positive because 1 and 0.5 are the lower bounds of  $x_1$  and  $x_2$ , respectively. Since the right-hand side must add up to an integer as well, this implies that

$$0.75 + 0.5(x_1 - 1) + 2.5(x_2 - 0.5) \geq 1. \quad (5.25)$$

Note, however, that this constraint is unsatisfied by  $\alpha$  since by construction all the elements on the left other than the fraction 0.75 are equal to zero under  $\alpha$ . This means that adding this constraint to the relaxed system will rule out this solution. On the other hand since it is implied by the integer system of constraints, it cannot remove any *integer* solution.

Let us generalize this example into a recipe for generating such cutting planes. The generalization refers also to the case of having variables assigned their upper bounds, and both negative and positive coefficients. In order to derive a Gomory cut from a constraint, the constraint has to satisfy two conditions: First, the assignment to the basic variable has to be fractional; Second, the assignments to all the nonbasic variables have to correspond to one of their bounds. The following recipe, which relies on these conditions, is based on a report by Dutertre and de Moura [71].

Consider the  $i$ -th constraint:

$$x_i = \sum_{x_j \in \mathcal{N}} a_{ij} x_j, \quad (5.26)$$

where  $x_i \in \mathcal{B}$ . Let  $\alpha$  be the assignment returned by the general simplex algorithm. Thus,

$$\alpha(x_i) = \sum_{x_j \in \mathcal{N}} a_{ij} \alpha(x_j). \quad (5.27)$$

We now partition the nonbasic variables to those that are currently assigned their lower bound and those that are currently assigned their upper bound:

$$\begin{aligned} J &= \{j \mid x_j \in \mathcal{N} \wedge \alpha(x_j) = l_j\} \\ K &= \{j \mid x_j \in \mathcal{N} \wedge \alpha(x_j) = u_j\}. \end{aligned} \quad (5.28)$$

Subtracting (5.27) from (5.26) taking the partition into account yields

$$x_i - \alpha(x_i) = \sum_{j \in J} a_{ij} (x_j - l_j) - \sum_{j \in K} a_{ij} (u_j - x_j). \quad (5.29)$$

Let  $f_0 = \alpha(x_i) - \lfloor \alpha(x_i) \rfloor$ . Since we assumed that  $\alpha(x_i)$  is not an integer then  $0 < f_0 < 1$ . We can now rewrite (5.29) as

$$x_i - \lfloor \alpha(x_i) \rfloor = f_0 + \sum_{j \in J} a_{ij} (x_j - l_j) - \sum_{j \in K} a_{ij} (u_j - x_j). \quad (5.30)$$

Note that the left-hand side is an integer. We now consider two cases.

- If  $\sum_{j \in J} a_{ij}(x_j - l_j) - \sum_{j \in K} a_{ij}(u_j - x_j) > 0$  then, since the right-hand side must be an integer,

$$f_0 + \sum_{j \in J} a_{ij}(x_j - l_j) - \sum_{j \in K} a_{ij}(u_j - x_j) \geq 1. \quad (5.31)$$

We now split  $J$  and  $K$  as follows:

$$\begin{aligned} J^+ &= \{j \mid j \in J \wedge a_{ij} > 0\} \\ J^- &= \{j \mid j \in J \wedge a_{ij} < 0\} \\ K^+ &= \{j \mid j \in K \wedge a_{ij} > 0\} \\ K^- &= \{j \mid j \in K \wedge a_{ij} < 0\} \end{aligned} \quad (5.32)$$

Gathering only the positive elements in the left-hand side of (5.31) gives us:

$$\sum_{j \in J^+} a_{ij}(x_j - l_j) - \sum_{j \in K^-} a_{ij}(u_j - x_j) \geq 1 - f_0, \quad (5.33)$$

or, equivalently,

$$\sum_{j \in J^+} \frac{a_{ij}}{1 - f_0}(x_j - l_j) - \sum_{j \in K^-} \frac{a_{ij}}{1 - f_0}(u_j - x_j) \geq 1. \quad (5.34)$$

- If  $\sum_{j \in J} a_{ij}(x_j - l_j) - \sum_{j \in K} a_{ij}(u_j - x_j) \leq 0$  then again, since the right-hand side must be an integer,

$$f_0 + \sum_{j \in J} a_{ij}(x_j - l_j) - \sum_{j \in K} a_{ij}(u_j - x_j) \leq 0. \quad (5.35)$$

Eq. (5.35) implies that

$$\sum_{j \in J^-} a_{ij}(x_j - l_j) - \sum_{j \in K^+} a_{ij}(u_j - x_j) \leq -f_0. \quad (5.36)$$

Dividing by  $-f_0$  gives us

$$- \sum_{j \in J^-} \frac{a_{ij}}{f_0}(x_j - l_j) + \sum_{j \in K^+} \frac{a_{ij}}{f_0}(u_j - x_j) \geq 1. \quad (5.37)$$

Note that the left-hand side of both (5.34) and (5.37) is greater than zero. Therefore these two equations imply

$$\begin{aligned} &\sum_{j \in J^+} \frac{a_{ij}}{1 - f_0}(x_j - l_j) - \sum_{j \in J^-} \frac{a_{ij}}{f_0}(x_j - l_j) \\ &+ \sum_{j \in K^+} \frac{a_{ij}}{f_0}(u_j - x_j) - \sum_{j \in K^-} \frac{a_{ij}}{1 - f_0}(u_j - x_j) \geq 1. \end{aligned} \quad (5.38)$$

Since each of the elements on the left-hand side is equal to zero under the current assignment  $\alpha$ , this assignment  $\alpha$  is ruled out by the new constraint. In other words, the solution to the linear problem augmented with the constraint is guaranteed to be different from the previous one.

## 5.4 Fourier–Motzkin Variable Elimination

### 5.4.1 Equality Constraints

Similarly to the simplex method, the Fourier–Motzkin variable elimination algorithm takes a conjunction of linear constraints over real variables. Let  $m$  denote the number of such constraints, and let  $x_1, \dots, x_n$  denote the variables used by these constraints.

As a first step, equality constraints of the following form are eliminated:

$$\sum_{j=1}^n a_{i,j} \cdot x_j = b_i . \quad (5.39)$$

We choose a variable  $x_j$  that has a nonzero coefficient  $a_{i,j}$  in an equality constraint  $i$ . Without loss of generality, we assume that  $x_n$  is the variable that is to be eliminated. The constraint (5.39) can be rewritten as

$$x_n = \frac{b_i}{a_{i,n}} - \sum_{j=1}^{n-1} \frac{a_{i,j}}{a_{i,n}} \cdot x_j . \quad (5.40)$$

Now we substitute the right-hand side of (5.40) for  $x_n$  into all the other constraints, and remove constraint  $i$ . This is iterated until all equalities are removed.

We are left with a system of inequalities of the form

$$\bigwedge_{i=1}^m \sum_{j=1}^n a_{i,j} x_j \leq b_i . \quad (5.41)$$

### 5.4.2 Variable Elimination

The basic idea of the variable elimination algorithm is to heuristically choose a variable and then to eliminate it by projecting its constraints onto the rest of the system, resulting in new constraints.


**Example 5.9.** Consider Fig. 5.3(a): the constraints

$$0 \leq x \leq 1, \quad 0 \leq y \leq 1, \quad \frac{3}{4} \leq z \leq 1 \quad (5.42)$$

form a cuboid. Projecting these constraints onto the  $x$  and  $y$  axes, and thereby eliminating  $z$ , results in a square which is given by the constraints

$$0 \leq x \leq 1, \quad 0 \leq y \leq 1 . \quad (5.43)$$

Figure 5.3(b) shows a triangle formed by the constraints



**Fig. 5.3.** Projection of constraints: (a) a cuboid is projected onto the  $x$  and  $y$  axes; (b) a triangle is projected onto the  $x$  axis

$$x \leq y + 10, \quad y \leq 15, \quad y \geq -x + 20 . \tag{5.44}$$

The projection of the triangle onto the  $x$  axis is a line given by the constraints

$$5 \leq x \leq 25 . \tag{5.45}$$

▀

Thus, the projection forms a new problem with one variable fewer, but possibly more constraints. This is done iteratively until all variables but one have been eliminated. The problem with one variable is trivially decidable.

The order in which the variables are eliminated may be predetermined, or adjusted dynamically to the current set of constraints. There are various heuristics for choosing the elimination order. A standard greedy heuristic gives priority to variables that produce fewer new constraints when eliminated.

Once again, assume that  $x_n$  is the variable chosen to be eliminated. The constraints are partitioned according to the coefficient of  $x_n$ . Consider the constraint with index  $i$ :

$$\sum_{j=1}^n a_{i,j} \cdot x_j \leq b_i . \tag{5.46}$$

By splitting the sum, (5.46) can be rewritten into

$$a_{i,n} \cdot x_n \leq b_i - \sum_{j=1}^{n-1} a_{i,j} \cdot x_j . \tag{5.47}$$

If  $a_{i,n}$  is zero, the constraint can be disregarded when we are eliminating  $x_n$ . Otherwise, we divide by  $a_{i,n}$ . If  $a_{i,n}$  is positive, we obtain

$$x_n \leq \frac{b_i}{a_{i,n}} - \sum_{j=1}^{n-1} \frac{a_{i,j}}{a_{i,n}} \cdot x_j . \tag{5.48}$$

$\beta_i$ 

Thus, if  $a_{i,n} > 0$ , the constraint is an *upper bound* on  $x_n$ . If  $a_{i,n} < 0$ , the constraint is a *lower bound*. We denote the right-hand side of (5.48) by  $\beta_i$ .

### Unbounded Variables

It is possible that a variable is not bounded both ways, i.e., it has either only upper bounds or only lower bounds. Such variables are called **unbounded variables**. Unbounded variables can be simply removed from the system together with all constraints that use them. Removing these constraints can make other variables unbounded. Thus, this simplification stage iterates until no such variables remain.

### Bounded Variables

If  $x_n$  has both an upper and a lower bound, the algorithm enumerates all pairs of lower and upper bounds. Let  $u \in \{1, \dots, m\}$  denote the index of an upper-bound constraint, and  $l \in \{1, \dots, m\}$  denote the index of a lower-bound constraint for  $x_n$ , where  $l \neq u$ . For each such pair, we have

$$\beta_l \leq x_n \leq \beta_u. \quad (5.49)$$

The following new constraint is added:

$$\beta_l \leq \beta_u. \quad (5.50)$$

The Formula (5.50) may simplify to  $0 \leq b_k$ , where  $b_k$  is some constant smaller than 0. In this case, the algorithm has found a *conflicting* pair of constraints and concludes that the problem is unsatisfiable. Otherwise, all constraints that involve  $x_n$  are removed. The new problem is solved recursively as before.

**Example 5.10.** Consider the following set of constraints:

$$\begin{aligned} x_1 - x_2 &\leq 0 \\ x_1 &\quad -x_3 \leq 0 \\ -x_1 + x_2 + 2x_3 &\leq 0 \\ &\quad -x_3 \leq -1. \end{aligned} \quad (5.51)$$

Suppose we decide to eliminate the variable  $x_1$  first. There are two upper bounds on  $x_1$ , namely  $x_1 \leq x_2$  and  $x_1 \leq x_3$ , and one lower bound, which is  $x_2 + 2x_3 \leq x_1$ .

Using  $x_1 \leq x_2$  as the upper bound, we obtain a new constraint  $2x_3 \leq 0$ , and using  $x_1 \leq x_3$  as the upper bound, we obtain a new constraint  $x_2 + x_3 \leq 0$ . Constraints involving  $x_1$  are removed from the problem, which results in the following new set:

$$\begin{aligned} 2x_3 &\leq 0 \\ x_2 + x_3 &\leq 0 \\ -x_3 &\leq -1. \end{aligned} \quad (5.52)$$

Next, observe that  $x_2$  is unbounded (as it has no lower bound), and hence the second constraint can be eliminated, which simplifies the formula. We therefore progress by eliminating  $x_2$  and all the constraints that contain it:

$$\begin{aligned} 2x_3 &\leq 0 \\ -x_3 &\leq -1 \end{aligned} \quad (5.53)$$

Only the variable  $x_3$  remains, with a lower and an upper bound. Combining the two into a new constraint results in  $1 \leq 0$ , which is a contradiction. Thus, the system is unsatisfiable.  $\blacksquare$

The simplex method in its basic form, as described in Sect. 5.2, allows only nonstrict ( $\leq$ ) inequalities.<sup>4</sup> The Fourier–Motzkin method, on the other hand, can easily be extended to handle a combination of strict ( $<$ ) and nonstrict inequalities: if either the lower or the upper bound is a strict inequality, then so is the resulting constraint.

### 5.4.3 Complexity

In each iteration, the number of constraints can increase in the worst case from  $m$  to  $m^2/4$ , which results overall in  $m^{2^n}/4^n$  constraints. Thus, Fourier–Motzkin variable elimination is only suitable for a relatively small set of constraints and a small number of variables.

## 5.5 The Omega Test

### 5.5.1 Problem Description

The Omega test is an algorithm to decide the satisfiability of a conjunction of linear constraints over integer variables. Each conjunct is assumed to be either an equality of the form

$$\sum_{i=1}^n a_i x_i = b \quad (5.54)$$

or a nonstrict inequality of the form

$$\sum_{i=1}^n a_i x_i \leq b. \quad (5.55)$$

The coefficients  $a_i$  are assumed to be integers; if they are not, by making use of the assumption that the coefficients are rational, the problem can be transformed into one with integer coefficients by multiplying the constraints

<sup>4</sup> There are extensions of Simplex to strict inequalities. See, for example, [70].

by the least common multiple of the denominators. In Sect. 5.6, we show how strict inequalities can be transformed into nonstrict inequalities.

The runtime of the Omega test depends on the size of the coefficients  $a_i$ . It is therefore desirable to transform the constraints such that small coefficients are obtained. This can be done by dividing the coefficients  $a_1, \dots, a_n$  of each constraint by their greatest common divisor  $g$ . The resulting constraint is called *normalized*. If the constraint is an equality constraint, this results in

$$\sum_{i=1}^n \frac{a_i}{g} x_i = \frac{b}{g}. \quad (5.56)$$

If  $g$  does not divide  $b$  exactly, the system is unsatisfiable. If the constraint is an inequality, one can tighten the constraint by rounding down the constant:

$$\sum_{i=1}^n \frac{a_i}{g} x_i \leq \left\lfloor \frac{b}{g} \right\rfloor. \quad (5.57)$$

More simplifications of this kind are described in Sect. 5.6.

**Example 5.11.** The equality  $3x + 3y = 2$  can be normalized to  $x + y = 2/3$ , which is unsatisfiable. The constraint  $8x + 6y \leq 0$  can be normalized to obtain  $4x + 3y \leq 0$ . The constraint  $1 \leq 4y$  can be tightened to obtain  $1 \leq y$ . ■

The Omega test is a variant of the Fourier–Motzkin variable elimination algorithm (Sect. 5.4). As in the case of that algorithm, equality and inequality constraints are treated separately; all equality constraints are removed before inequalities are considered.

### 5.5.2 Equality Constraints

In order to eliminate an equality of the form of (5.54), we first check if there is a variable  $x_j$  with a coefficient 1 or  $-1$ , i.e.,  $|a_j| = 1$ . If yes, we transform the constraint as follows. Without loss of generality, assume  $j = n$ . We isolate  $x_n$ :

$$x_n = \frac{b}{a_n} - \sum_{i=1}^{n-1} \frac{a_i}{a_n} x_i. \quad (5.58)$$

The variable  $x_n$  can now be substituted by the right-hand side of (5.58) in all constraints.

If there is no variable with a coefficient 1 or  $-1$ , we cannot simply divide by the coefficient, as this would result in nonintegral coefficients. Instead, the algorithm proceeds as follows: it determines the variable that has the nonzero coefficient with the smallest absolute value. Assume again that  $x_n$  is chosen, and that  $a_n > 0$ . The Omega test transforms the constraints iteratively until some coefficient becomes 1 or  $-1$ . The variable with that coefficient can then be eliminated as above.

For this transformation, a new binary operator  $\widehat{\text{mod}}$ , called **symmetric modulo**, is defined as follows:

$$\widehat{a \bmod b}$$

$$a \widehat{\text{mod}} b \doteq a - b \cdot \left\lfloor \frac{a}{b} + \frac{1}{2} \right\rfloor . \quad (5.59)$$

The symmetric modulo operator is very similar to the usual modular arithmetic operator. If  $a \bmod b < b/2$ , then  $a \widehat{\text{mod}} b = a \bmod b$ . If  $a \bmod b$  is greater than or equal to  $b/2$ ,  $b$  is deducted, and thus

$$a \widehat{\text{mod}} b = \begin{cases} a \bmod b & : a \bmod b < b/2 \\ (a \bmod b) - b & : \text{otherwise} . \end{cases} \quad (5.60)$$

We leave the proof of this equivalence as an exercise (see Problem 5.12).

Our goal is to derive a term that can replace  $x_n$ . For this purpose, we define  $m \doteq a_n + 1$ , introduce a new variable  $\sigma$ , and add the following new constraint:

$$\sum_{i=1}^n (a_i \widehat{\text{mod}} m) x_i = m\sigma + b \widehat{\text{mod}} m . \quad (5.61)$$

We split the sum on the left-hand side to obtain

$$(a_n \widehat{\text{mod}} m) x_n = m\sigma + b \widehat{\text{mod}} m - \sum_{i=1}^{n-1} (a_i \widehat{\text{mod}} m) x_i . \quad (5.62)$$

Since  $a_n \widehat{\text{mod}} m = -1$  (see Problem 5.14), this simplifies to:

$$x_n = -m\sigma - b \widehat{\text{mod}} m + \sum_{i=1}^{n-1} (a_i \widehat{\text{mod}} m) x_i . \quad (5.63)$$

The right-hand side of (5.63) is used to replace  $x_n$  in all constraints. Any equality from the original problem (5.54) is changed as follows:

$$\sum_{i=1}^{n-1} a_i x_i + a_n \left( -m\sigma - b \widehat{\text{mod}} m + \sum_{i=1}^{n-1} (a_i \widehat{\text{mod}} m) x_i \right) = b , \quad (5.64)$$

which can be rewritten as

$$-a_n m\sigma + \sum_{i=1}^{n-1} (a_i + a_n (a_i \widehat{\text{mod}} m)) x_i = b + a_n (b \widehat{\text{mod}} m) . \quad (5.65)$$

Since  $a_n = m - 1$ , this simplifies to

$$-a_n m\sigma + \sum_{i=1}^{n-1} ((a_i - (a_i \widehat{\text{mod}} m)) + m(a_i \widehat{\text{mod}} m)) x_i = b - (b \widehat{\text{mod}} m) + m(b \widehat{\text{mod}} m) . \quad (5.66)$$



Note that  $a_i - (a_i \widehat{\text{mod}} m)$  is equal to  $m\lfloor a_i/m + 1/2 \rfloor$ , and thus all terms are divisible by  $m$ . Dividing (5.66) by  $m$  results in

$$-a_n\sigma + \sum_{i=1}^{n-1} (\lfloor a_i/m + 1/2 \rfloor + (a_i \widehat{\text{mod}} m))x_i = \lfloor b/m + 1/2 \rfloor + (b \widehat{\text{mod}} m). \quad (5.67)$$

The absolute value of the coefficient of  $\sigma$  is the same as the absolute value of the original coefficient  $a_n$ , and it seems that nothing has been gained by this substitution. However, observe that the coefficient of  $x_i$  can be bounded as follows (see Problem 5.13):

$$|\lfloor a_i/m + 1/2 \rfloor + (a_i \widehat{\text{mod}} m)| \leq \frac{5}{6}|a_i|. \quad (5.68)$$

Thus, the absolute values of the coefficients in the equality are strictly smaller than their previous values. As the coefficients are always integral, repeated application of equality elimination eventually generates a coefficient of 1 or  $-1$  on some variable. This variable can then be eliminated directly, as described earlier (see (5.58)).

**Example 5.12.** Consider the following formula:

$$\begin{aligned} -3x_1 + 2x_2 &= 0 \\ 3x_1 + 4x_2 &= 3. \end{aligned} \quad (5.69)$$

The variable  $x_2$  has the coefficient with the smallest absolute value ( $a_2 = 2$ ). Thus,  $m = a_2 + 1 = 3$ , and we add the following constraint (see (5.61)):

$$(-3 \widehat{\text{mod}} 3)x_1 + (2 \widehat{\text{mod}} 3)x_2 = 3\sigma. \quad (5.70)$$

This simplifies to  $x_2 = -3\sigma$ . Substituting  $-3\sigma$  for  $x_2$  results in the following problem:

$$\begin{aligned} -3x_1 - 6\sigma &= 0 \\ 3x_1 - 12\sigma &= 3. \end{aligned} \quad (5.71)$$

Division by  $m$  results in

$$\begin{aligned} -x_1 - 2\sigma &= 0 \\ x_1 - 4\sigma &= 1. \end{aligned} \quad (5.72)$$

As expected, the coefficient of  $x_1$  has decreased. We can now substitute  $x_1$  by  $4\sigma + 1$ , and obtain  $-6\sigma = 1$ , which is unsatisfiable.  $\blacksquare$

### 5.5.3 Inequality Constraints

Once all equalities have been eliminated, the algorithm attempts to find a solution for the remaining inequalities. The control flow of Algorithm 5.5.1 is illustrated in Fig. 5.4. As in the Fourier–Motzkin procedure, the first step is to choose a variable to be eliminated. Subsequently, the three subprocedures

*Real-Shadow*, *Dark-Shadow*, and *Gray-Shadow* produce new constraint sets, which are solved recursively.

Note that many of the subproblems generated by the recursion are actually identical. An efficient implementation uses a hash table that stores the solutions of previously solved problems.

**Algorithm 5.5.1: OMEGA-TEST**

**Input:** A conjunction of constraints  $C$

**Output:** “Satisfiable” if  $C$  is satisfiable, and “Unsatisfiable” otherwise

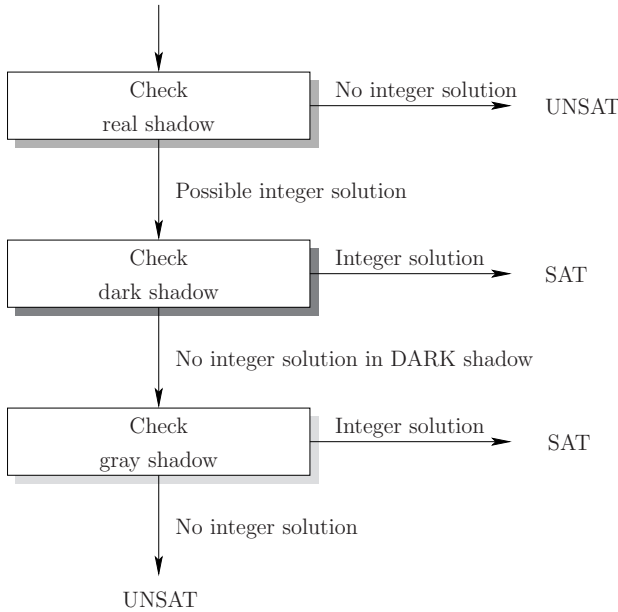
```

1. if  $C$  only contains one variable then
2.   Solve and return result;           ▷ (solving this problem is trivial)
3.
4. Otherwise, choose a variable  $v$  that occurs in  $C$ ;
5.  $C_R := \text{Real-Shadow}(C, v)$ ;
6. if OMEGA-TEST( $C_R$ ) = “Unsatisfiable” then           ▷ Recursive call
7.   return “Unsatisfiable”;
8.
9.  $C_D := \text{Dark-Shadow}(C, v)$ ;
10. if OMEGA-TEST( $C_D$ ) = “Satisfiable” then           ▷ Recursive call
11.   return “Satisfiable”;
12.
13. if  $C_R = C_D$  then                                   ▷ Exact projection?
14.   return “Unsatisfiable”;
15.
16.  $C_G^1, \dots, C_G^n := \text{Gray-Shadow}(C, v)$ ;
17. for all  $i \in \{1, \dots, n\}$  do
18.   if OMEGA-TEST( $C_G^i$ ) = “Satisfiable” then           ▷ Recursive call
19.     return “Satisfiable”;
20.
21. return “Unsatisfiable”;

```

### Checking the Real Shadow

Even though the Omega test is concerned with constraints over integers, the first step is to check if there are integer solutions in the relaxed problem, which is called the *real shadow*. The real shadow is the same projection that the Fourier–Motzkin procedure uses. The Omega test is then called recursively to check if the projection contains an integer. If there is no such integer, then there is no integer solution to the original system either, and the algorithm concludes that the system is unsatisfiable.



**Fig. 5.4.** Overview of the Omega test

Assume that the variable to be eliminated is denoted by  $z$ . As in the case of the Fourier–Motzkin procedure, all pairs of lower and upper bounds have to be considered. Variables that are not bounded both ways can be removed, together with all constraints that contain them.

Let  $\beta \leq bz$  and  $cz \leq \gamma$  be constraints, where  $c$  and  $b$  are positive integer constants and  $\gamma$  and  $\beta$  denote the remaining linear expressions. Consequently,  $\beta/b$  is a lower bound on  $z$ , and  $\gamma/c$  is an upper bound on  $z$ . The new constraint is obtained by multiplying the lower bound by  $c$  and the upper bound by  $b$ :

$$\begin{array}{ccc}
 \text{Lower bound} & & \text{Upper bound} \\
 \hline
 \beta \leq bz & & cz \leq \gamma \\
 c\beta \leq cbz & & cbz \leq b\gamma
 \end{array} \tag{5.73}$$

The existence of such a variable  $z$  implies

$$c\beta \leq b\gamma. \tag{5.74}$$

**Example 5.13.** Consider the following set of constraints:

$$\begin{array}{l}
 2y \leq x \\
 8y \geq 2 + x \\
 2y \leq 3 - x
 \end{array} \tag{5.75}$$

The triangle spanned by these constraints is depicted in Fig. 5.5. Assume that we decide to eliminate  $x$ . In this case, the combination of the two constraints

$2y \leq x$  and  $8y \geq 2 + x$  results in  $8y - 2 \geq 2y$ , which simplifies to  $y \geq 1/3$ . The two constraints  $2y \leq x$  and  $2y \leq 3 - x$  combine into  $2y \leq 3 - 2y$ , which simplifies to  $y \leq 3/4$ . Thus,  $1/3 \leq y \leq 3/4$  must hold, which has no integer solution. The set of constraints is therefore unsatisfiable.  $\blacksquare$




Fig. 5.5. Computing the real shadow: eliminating  $x$

The converse of this observation does not hold, i.e., if we find an integer solution within the real shadow, this does not guarantee that the original set of constraints has an integer solution. This is illustrated by the following example.




Fig. 5.6. Computing the real shadow: eliminating  $y$

**Example 5.14.** Consider the same set of constraints as in Example 5.13. This time, eliminate  $y$  instead of  $x$ . This projection is depicted in Fig. 5.6.

We obtain  $2/3 \leq x \leq 2$ , which has two integer solutions. The triangle, on the other hand, contains no integer solution.  $\blacksquare$

The real shadow is an overapproximating projection, as it contains more solutions than does the original problem. The next step in the Omega test is to compute an underapproximating projection, i.e., if that projection contains an integer solution, so does the original problem. This projection is called the *dark shadow*.

### Checking the Dark Shadow

The name *dark shadow* is motivated by optics. Assume that the object we are projecting is partially translucent. Places that are “thicker” will project a darker shadow. In particular, a dark area in the shadow where the object is thicker than 1 must have at least one integer above it.


After the first phase of the algorithm, we know that there is a solution to the real shadow, i.e.,  $c\beta \leq b\gamma$ . We now aim at determining if there is an integer  $z$  such that  $c\beta \leq cz \leq b\gamma$ , which is equivalent to

$$\exists z \in \mathbb{Z}. \frac{\beta}{b} \leq z \leq \frac{\gamma}{c}. \quad (5.76)$$

Assume that (5.76) does not hold. Let  $i$  denote  $\lfloor \beta/b \rfloor$ , i.e., the largest integer that is smaller than  $\beta/b$ . Since we have assumed that there is no integer between  $\beta/b$  and  $\gamma/c$ ,

$$i < \frac{\beta}{b} \leq \frac{\gamma}{c} < i + 1 \quad (5.77)$$

holds. This situation is illustrated in Fig. 5.7.



**Fig. 5.7.** Computing the dark shadow

Since  $\beta/b$  and  $\gamma/c$  are not integers themselves, the distances from these points to the closest integer are greater than the fractions  $1/b$  and  $1/c$ , respectively:

$$\frac{\beta}{b} - i \geq \frac{1}{b} \quad (5.78)$$

$$i + 1 - \frac{\gamma}{c} \geq \frac{1}{c}. \quad (5.79)$$

The proof is left as an exercise (Problem 5.11). By summing (5.78) and (5.79), we obtain

$$\frac{\beta}{b} + 1 - \frac{\gamma}{c} \geq \frac{1}{c} + \frac{1}{b}, \quad (5.80)$$

which is equivalent to

$$c\beta - b\gamma \geq -cb + c + b. \quad (5.81)$$

By multiplying this inequality by  $-1$ , we obtain

$$b\gamma - c\beta \leq cb - c - b. \quad (5.82)$$

In order to show a contradiction to our assumption, we need to show the negation of (5.82). Exploiting the fact that  $c, b$  are integers, the negation of (5.82) is

$$b\gamma - c\beta \geq cb - c - b + 1, \quad (5.83)$$

or simply

$$b\gamma - c\beta \geq (c-1)(b-1). \quad (5.84)$$

Thus, if (5.84) holds, our assumption is wrong, which means that we have a guarantee that there exists an integer solution.

Observe that if either  $c = 1$  or  $b = 1$ , the formula (5.84) is identical to the real shadow (5.74), i.e., the dark and real shadow are the same. In this case, the projection is exact, and it is sufficient to check the *real shadow*. When choosing variables to eliminate, preference should be given to variables that result in an exact projection, that is, to variables with coefficient 1.

### Checking the Gray Shadow

We know that any integer solution must also be in the real shadow. Let **R** denote this area. Now assume that we have found no integer in the dark shadow. Let **D** denote the area of the dark shadow.

Thus, if **R** and **D** do not coincide, there is only one remaining area in which an integer solution can be found: an area around the dark shadow, which, staying within the optical analogy, is called the *gray shadow*.

Any solution must satisfy

$$c\beta \leq cbz \leq b\gamma. \quad (5.85)$$

Furthermore, we already know that the dark shadow does not contain an integer, and thus we can exclude this area from the search. Therefore, besides (5.85), any solution has to satisfy (5.82):

**R**

**D**

$$c\beta \leq cbz \leq b\gamma \quad \wedge \quad b\gamma - c\beta \leq cb - c - b. \quad (5.86)$$

This is equivalent to

$$c\beta \leq cbz \leq b\gamma \quad \wedge \quad b\gamma \leq cb - c - b + c\beta, \quad (5.87)$$

which implies

$$c\beta \leq cbz \leq cb - c - b + c\beta. \quad (5.88)$$

Dividing by  $c$ , we obtain

$$\beta \leq bz \leq \beta + \frac{cb - c - b}{c}. \quad (5.89)$$

The Omega test proceeds by simply trying possible values of  $bz$  between these two bounds. Thus, a new constraint

$$bz = \beta + i \quad (5.90)$$

is formed and combined with the original problem for each integer  $i$  in the range  $0, \dots, (cb - c - b)/c$ . If any one of the resulting new problems has a solution, so does the original problem.

The number of subproblems can be reduced by determining the largest coefficient  $c$  of  $z$  in any upper bound for  $z$ . The new constraints generated for the other upper bounds are already covered by the constraints generated for the upper bound with the largest  $c$ .

## 5.6 Preprocessing

In this section, we examine several simple preprocessing steps for both linear and integer linear systems without objective functions. Preprocessing the set of constraints can be done regardless of the decision procedure chosen.

### 5.6.1 Preprocessing of Linear Systems

Two simple preprocessing steps for linear systems are the following:

1. Consider the set of constraints

$$x_1 + x_2 \leq 2, \quad x_1 \leq 1, \quad x_2 \leq 1. \quad (5.91)$$

The first constraint is redundant. In general, for a set:

$$S = \left\{ a_0x_0 + \sum_{j=1}^n a_jx_j \leq b, \quad l_j \leq x_j \leq u_j \text{ for } j = 0, \dots, n \right\}, \quad (5.92)$$

the constraint

$$a_0x_0 + \sum_{j=1}^n a_jx_j \leq b \quad (5.93)$$

is redundant if

$$\sum_{j|a_j>0} a_ju_j + \sum_{j|a_j<0} a_jl_j \leq b. \quad (5.94)$$

To put this in words, a “ $\leq$ ” constraint in the above form is redundant if assigning values equal to their upper bounds to all of its variables that have a positive coefficient, and assigning values equal to their lower bounds to all of its variables that have a negative coefficient, results in a value less than or equal to  $b$ , the constant on the right-hand side of the inequality.

2. Consider the following set of constraints:

$$2x_1 + x_2 \leq 2, \quad x_2 \geq 4, \quad x_1 \leq 3. \quad (5.95)$$

From the first and second constraints,  $x_1 \leq -1$  can be derived, which means that the bound  $x_1 \leq 3$  can be tightened. In general, if  $a_0 > 0$ , then

$$x_0 \leq \left( b - \sum_{j|j>0, a_j>0} a_jl_j - \sum_{j|a_j<0} a_ju_j \right) / a_0, \quad (5.96)$$

and if  $a_0 < 0$ , then

$$x_0 \geq \left( b - \sum_{j|a_j>0} a_jl_j - \sum_{j|j>0, a_j<0} a_ju_j \right) / a_0. \quad (5.97)$$

### 5.6.2 Preprocessing of Integer Linear Systems

The following preprocessing steps are applicable to integer linear systems:

1. Multiply every constraint by the smallest common multiple of the coefficients and constants in this constraint, in order to obtain a system with integer coefficients.<sup>5</sup>
2. After the previous preprocessing has been applied, strict inequalities can be transformed into nonstrict inequalities as follows:

$$\sum_{1 \leq i \leq n} a_ix_i < b \quad (5.98)$$

is replaced with

$$\sum_{1 \leq i \leq n} a_ix_i \leq b - 1. \quad (5.99)$$

The case in which  $b$  is fractional is handled by the previous preprocessing step.

<sup>5</sup> This assumes that the coefficients and constants in the system are rational. The case in which the coefficients can be nonrational is of little value and is rarely considered in the literature.



For the special case of **0–1 linear systems** (integer linear systems in which all the variables are constrained to be either 0 or 1), some preprocessing steps are illustrated by the following examples:

1. Consider the constraint

$$5x_1 - 3x_2 \leq 4, \quad (5.100)$$

from which we can conclude that

$$x_1 = 1 \implies x_2 = 1. \quad (5.101)$$

Hence, the constraint

$$x_1 \leq x_2 \quad (5.102)$$

can be added.

2. From

$$x_1 + x_2 \leq 1, \quad x_2 \geq 1, \quad (5.103)$$

we can conclude  $x_1 = 0$ .

Generalization of these examples is left for Problem 5.8.

## 5.7 Difference Logic

### 5.7.1 Introduction

A popular fragment of linear arithmetic is called **difference logic**.

**Definition 5.15 (difference logic).** *The syntax of a formula in difference logic is defined by the following rules:*

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid \text{atom} \\ \text{atom} &: \text{identifier} - \text{identifier} \text{ op constant} \\ \text{op} &: \leq \mid < \end{aligned}$$

Here, we consider the case in which the variables are defined over  $\mathbb{Q}$ , the rationals. A similar definition exists for the case in which the variables are defined over  $\mathbb{Z}$  (see Problem 5.18). Solving both variants is polynomial, whereas, recall, linear arithmetic over  $\mathbb{Z}$  is NP-complete.

Some other convenient operands can be modeled with the grammar above:

- $x - y = c$  is the same as  $x - y \leq c \wedge y - x \leq -c$ .
- $x - y \geq c$  is the same as  $y - x \leq -c$ .
- $x - y > c$  is the same as  $y - x < -c$ .
- A constraint with one variable such as  $x < 5$  can be rewritten as  $x - x_0 < 5$ , where  $x_0$  is a special variable not used so far in the formula, called the “zero variable”. In any satisfying assignment, its value must be 0.

As an example,

$$x < y + 5 \wedge y \leq 4 \wedge x = z - 1 \quad (5.104)$$

can be rewritten in difference logic as

$$x - y < 5 \wedge y - x_0 \leq 4 \wedge x - z \leq -1 \wedge z - x \leq 1. \quad (5.105)$$

A more important variant, however, is one in which an arbitrary Boolean structure is permitted. We describe one application of this variant by the following example.

**Example 5.16.** We are given a finite set of  $n$  jobs, each of which consists of a chain of operations. There is a finite set of  $m$  machines, each of which can handle at most one operation at a time. Each operation needs to be performed during an uninterrupted period of given length on a given machine. The **job-shop scheduling** problem is to find a schedule, that is, an allocation of the operations to time intervals on the machines that has a minimal total length.

More formally, given a set of machines

$$M = \{m_1, \dots, m_m\}, \quad (5.106)$$

job  $J^i$  with  $i \in \{1, \dots, n\}$  is a sequence of  $n_i$  pairs of the form (machine, duration):

$$J^i = (m_1^i, d_1^i), \dots, (m_{n_i}^i, d_{n_i}^i), \quad (5.107)$$

such that  $m_1^i, \dots, m_{n_i}^i$  are elements of  $M$ . The durations can be assumed to be rational numbers. We denote by  $O$  the multiset of all operations from all jobs. For an operation  $v \in O$ , we denote its machine by  $M(v)$  and its duration by  $\tau(v)$ .

A schedule is a function that defines, for each operation  $v$ , its starting time  $S(v)$  on its specified machine  $M(v)$ . A schedule  $S$  is *feasible* if the following three constraints hold.

First, the starting time of all operations is greater than or equal to 0:

$$\forall v \in O. S(v) \geq 0. \quad (5.108)$$

Second, for every pair of consecutive operations  $v_i, v_j \in O$  in the same job, the second operation does not start before the first ends:

$$S(v_i) + \tau(v_i) \leq S(v_j). \quad (5.109)$$

Finally, every pair of different operations  $v_i, v_j \in O$  scheduled on the same machine ( $M(v_i) = M(v_j)$ ) is mutually exclusive:

$$S(v_i) + \tau(v_i) \leq S(v_j) \vee S(v_j) + \tau(v_j) \leq S(v_i). \quad (5.110)$$

The length of the schedule  $S$  is defined as

$$\max_{v \in O} S(v) + \tau(v), \quad (5.111)$$

and the objective is to find a schedule  $S$  that minimizes this length. As usual, we can define the decision problem associated with this optimization problem by removing the objective function and adding a constraint that forces the value of this function to be smaller than some constant.

It should be clear that a job-shop scheduling problem can be formulated with difference logic. Note the disjunction in (5.110).  $\blacksquare$

### 5.7.2 A Decision Procedure for Difference Logic

Recall that in this chapter we present only decision procedures for conjunctive fragments, and postpone the problem of solving the general case to Chap. 11.

**Definition 5.17 (inequality graph for nonstrict inequalities).** *Let  $S$  be a set of difference predicates and let the inequality graph  $G(V, E)$  be the graph comprising of one edge  $(x, y)$  with weight  $c$  for every constraint of the form  $x - y \leq c$  in  $S$ .*

Given a difference logic formula  $\varphi$  with nonstrict inequalities only, the inequality graph corresponding to the set of difference predicates in  $\varphi$  can be used for deciding  $\varphi$ , on the basis of the following theorem.

**Theorem 5.18.** *Let  $\varphi$  be a conjunction of difference constraints, and let  $G$  be the corresponding inequality graph. Then  $\varphi$  is satisfiable if and only if there is no negative cycle in  $G$ .*

The proof of this theorem is left as an exercise (Problem 5.15). The extension of Definition 5.17 and Theorem 5.18 to general difference logic (which includes both strict and nonstrict inequalities) is left as an exercise as well (see Problem 5.16).

By Theorem 5.18, deciding a difference logic formula amounts to searching for a negative cycle in a graph. This can be done with the **Bellman–Ford algorithm** [54] for finding the single-source shortest paths in a directed weighted graph, in time  $O(|V| \cdot |E|)$  (to make the graph single-source, we introduce a new node and add an edge with weight 0 from this node to each of the roots of the original graph). Although finding the shortest paths is not our goal, we exploit a side-effect of this algorithm: if there exists a negative cycle in the graph, the algorithm finds it and aborts.

## 5.8 Problems

### 5.8.1 Warm-up Exercises

**Problem 5.1 (linear systems).** Consider the following linear system, which we denote by  $S$ :

$$\begin{aligned} x_1 &\geq -x_2 + \frac{11}{5} \\ x_1 &\leq x_2 + \frac{1}{2} \\ x_1 &\geq 3x_2 - 3 \end{aligned} \quad (5.112)$$

- (a) Check with simplex whether  $S$  is satisfiable, as described in Sect. 5.2.
- (b) Using the Fourier–Motzkin procedure, compute the range within which  $x_2$  has to lie in a satisfying assignment.
- (c) Consider a problem  $S'$ , similar to  $S$ , but where the variables are forced to be integer. Check with Branch and Bound whether  $S'$  is satisfiable. To solve the relaxed problem, you can use a simplex implementation (there are many of these on the Web).

### 5.8.2 The Simplex Method

**Problem 5.2 (simplex).** Compute a satisfying assignment for the following problem using the general simplex method:

$$\begin{aligned}
 2x_1 + 2x_2 + 2x_3 + 2x_4 &\leq 2 \\
 4x_1 + x_2 + x_3 - 4x_4 &\leq -2 \\
 x_1 + 2x_2 + 4x_3 + 2x_4 &= 4 \ .
 \end{aligned}
 \tag{5.113}$$

**Problem 5.3 (complexity).** Give a conjunction of linear constraints over reals with  $n$  variables (that is, the size of the instance is parameterized) such that the number of iterations of the general simplex algorithm is exponential in  $n$ .

**Problem 5.4 (difference logic with simplex).** What is the worst-case run time of the general simplex algorithm if applied to a conjunction of difference logic constraints?

**Problem 5.5 (strict inequalities with simplex).** Extend the general simplex algorithm with strict inequalities.

**Problem 5.6 (soundness).** Assume that the general simplex algorithm returns “UNSAT”. Show a method for deriving a proof of unsatisfiability.

### 5.8.3 Integer Linear Systems

**Problem 5.7 (complexity of ILP-feasibility).** Prove that the feasibility problem for integer linear programming is NP-hard.<sup>6</sup>

**Problem 5.8 (0–1 ILP).** A 0–1 integer linear system is an integer linear system in which all variables are constrained to be either 0 or 1. Show how a 0–1 integer linear system can be translated to a Boolean formula. What is the complexity of the translation?

<sup>6</sup> In fact it is NP-complete, but membership in NP is more difficult to prove. The proof makes use of a small-model-property argument.

**Problem 5.9 (simplifications for 0–1 ILP).** Generalize the simplification demonstrated in (5.100)–(5.103).

**Problem 5.10 (Gomory cuts).** Find Gomory cuts corresponding to the following results from the general simplex algorithm:

1.  $x_4 = x_1 - 2.5x_2 + 2x_3$  where  $\alpha := \{x_4 \mapsto 3.25, x_1 \mapsto 1, x_2 \mapsto -0.5, x_3 \mapsto 0.5\}$ ,  $x_2$  and  $x_3$  are at their upper bound and  $x_1$  is at its lower bound.
2.  $x_4 = -0.5x_1 - 2x_2 - 3.5x_3$  where  $\alpha := \{x_4 \mapsto 0.25, x_1 \mapsto 1, x_2 \mapsto 0.5, x_3 \mapsto 0.5\}$ ,  $x_1$  and  $x_3$  are at their lower bound and  $x_2$  is at its upper bound.

### 5.8.4 Omega Test

**Problem 5.11 (integer fractions).** Show that

$$i + 1 - \frac{\gamma}{c} \geq \frac{1}{c}.$$

**Problem 5.12 (eliminating equalities).** Show that

$$a \widehat{\text{mod}} b = \begin{cases} a \bmod b & : a \bmod b < b/2 \\ (a \bmod b) - b & : \text{otherwise} \end{cases} \quad (5.114)$$

holds. Use the fact that

$$a/b = \lfloor a/b \rfloor + \frac{a \bmod b}{b}.$$

**Problem 5.13 (eliminating equalities).** Show that the absolute values of the coefficients of the variables  $x_i$  are reduced to at most  $5/6$  of their previous values after substituting  $\sigma$ :

$$|\lfloor a_i/m + 1/2 \rfloor + (a_i \widehat{\text{mod}} m)| \leq 5/6 |a_i|. \quad (5.115)$$

**Problem 5.14 (eliminating equalities).** The elimination of  $x_n$  relies on the fact that the coefficient of  $x_n$  in the newly added constraint is  $-1$ . Let  $a_n$  denote the coefficient of  $x_n$  in the original constraint. Let  $m = a_n + 1$ , and assume that  $a_n \geq 2$ . Show that  $a_n \widehat{\text{mod}} m = -1$ .

### 5.8.5 Difference Logic

**Problem 5.15 (difference logic).** Prove Theorem 5.18.

**Problem 5.16 (inequality graphs for difference logic).** Extend Definition 5.17 and Theorem 5.18 to general difference logic formulas (i.e., where both strong and weak inequalities are allowed).

**Problem 5.17 (difference logic).** Give a reduction of difference logic to SAT. What is the complexity of the reduction?

**Problem 5.18 (integer difference logic).** Show a reduction from the problem of integer difference logic to difference logic.

## 5.9 Bibliographic Notes

The Fourier–Motzkin variable elimination algorithm is the earliest documented method for solving linear inequalities. It was discovered in 1826 by Fourier, and rediscovered by Motzkin in 1936.

The simplex method was introduced by Danzig in 1947 [55]. There are several variations of and improvements on this method, most notably the *revised simplex method*, which most industrial implementations use. This variant has an apparent advantage on large and sparse LP problems, which seem to characterize LP problems in practice. The variant of the general simplex algorithm that we presented in Sect. 5.2 was proposed by Dutertre and de Moura [70] in the context of DPLL( $T$ ), a technique we describe in Chap. 11. Its main advantage is that it works efficiently with incremental operations, i.e., constraints can be added and removed with little effort.

Linear programs are a very popular modeling formalism for solving a wide range of problems in science and engineering, finance, logistics and so on. See, for example, how LP is used for computing an optimal placement of gates in an integrated circuit [100]. The popularity of this method led to a large industry of LP solvers, some of which are sold for tens of thousands of dollars per copy. A classical reference to linear and integer linear programming is the book by Schrijver [174]. Other resources on the subject that we found useful include publications by Wolsey [204], Hillier and Lieberman [92], and Vanderbei [196].

Gomory cutting-planes are due to a paper published by Ralph Gomory in 1963 [89]. For many years, the operations research community considered Gomory cuts impractical for large problems. There were several refinements of the original method and empirical studies that revived this technique, especially in the context of the related optimization problem. See, for example, the work of Balas et al. [72]. The variant we described is suitable for working with the general simplex algorithm and its description here is based on [71].

The Omega test was introduced by Pugh as a method for deciding integer linear arithmetic within an optimizing compiler [160]. It is an extension of the

Fourier–Motzkin variable elimination. For an example of an application of the Omega test inside a Fortran compiler, see [2]. A much earlier work following similar lines to those of the omega test is by Paul Williams [199]. Williams’ work, in turn, is inspired by Presburger’s paper from 1929 [159].

Difference logic was recognized as an interesting fragment of linear arithmetic by Pratt [158]. He considered “separation theory”, which is the conjunctive fragment of what we call difference logic. He observed that most inequalities in verification conditions are of this form. Disjunctive difference logic was studied in M. Mahfoudh’s PhD thesis [119] and in [120], among other places. A reduction of difference logic to SAT was studied in [187] (in this particular paper and some later papers, this theory fragment is called “separation logic”, after Pratt’s separation theory – not to be confused with the separation logic that is discussed in Chap. 8). The main reason for the renewed interest in this fragment is due to interest in **timed automata**: the verification conditions arising in this problem domain are difference logic formulas.

In general, the amount of research and writing on linear systems is immense, and in fact most universities offer courses dedicated to this subject. Most of the research was and still is conducted in the operations research community.

## 5.10 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$l_i, u_i$	Constants bounding the $i$ -th variable from below and above	113
$m$	The number of linear constraints in the original problem formulation	114
$n$	The number of variables in the original problem formulation	114
$A$	Coefficient matrix	115
$\mathbf{x}$	The vector of the variables in the original problem formulation	115
$\mathcal{B}, N$	The sets of basic and nonbasic variables, respectively	116
<i>continued on next page</i>		

*continued from previous page*

<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$\alpha$	A full assignment (to both basic and nonbasic variables)	116
$\theta$	See (5.13)	118
$\beta_i$	Upper or lower bound	128
$\widehat{\text{mod}}$	Symmetric modulo	131



## Bit Vectors

---

### 6.1 Bit-Vector Arithmetic

The design of computer systems is error-prone, and, thus, decision procedures for reasoning about such systems are highly desirable. A computer system uses *bit vectors* to encode information, for example numbers. Owing to the finite domain of these bitvectors, the semantics of operations such as addition no longer matches what we are used to when reasoning about unbounded types, for example the natural numbers.

#### 6.1.1 Syntax

The subset of bit-vector arithmetic that we consider is defined by the following grammar:

$$\begin{aligned}
 \textit{formula} & : \textit{formula} \wedge \textit{formula} \mid \neg \textit{formula} \mid (\textit{formula}) \mid \textit{atom} \\
 \textit{atom} & : \textit{term} \textit{rel} \textit{term} \mid \textit{Boolean-Identifier} \mid \textit{term}[\textit{constant}] \\
 \textit{rel} & : < \mid = \\
 \textit{term} & : \textit{term} \textit{op} \textit{term} \mid \textit{identifier} \mid \sim \textit{term} \mid \textit{constant} \mid \textit{atom} ? \textit{term} : \textit{term} \mid \\
 & \quad \textit{term}[\textit{constant} : \textit{constant}] \mid \textit{ext}(\textit{term}) \\
 \textit{op} & : + \mid - \mid \cdot \mid / \mid \ll \mid \gg \mid \& \mid \mid \mid \oplus \mid \circ
 \end{aligned}$$

As usual, other useful operators such as “ $\vee$ ”, “ $\neq$ ”, and “ $\geq$ ” can be obtained using Boolean combinations of the operators that appear in the grammar. Most operators have a straightforward meaning, but a few operators are unique to bit-vector arithmetic. The unary operator “ $\sim$ ” denotes bitwise negation. The function *ext* denotes sign and zero extension (the meanings of these operators are explained in Sect. 6.1.3). The ternary operator *c?a:b* is a case-split: the operator evaluates to *a* if *c* holds, and to *b* otherwise. The operators “ $\ll$ ” and “ $\gg$ ” denote left and right shifts, respectively. The operator “ $\oplus$ ” denotes bitwise XOR. The binary operator “ $\circ$ ” denotes concatenation of bit vectors.

## Motivation

As an example to describe our motivation, the following formula obviously holds over the integers:

$$(x - y > 0) \iff (x > y). \quad (6.1)$$

If  $x$  and  $y$  are interpreted as finite-width bit vectors, however, this equivalence no longer holds, owing to possible **overflow** of the subtraction operation. As another example, consider the following small C program:

```
unsigned char number = 200;
number = number + 100;
printf("Sum: %d\n", number);
```

This program may return a surprising result, as most architectures use eight bits to represent variables with type `unsigned char`:

$$\begin{array}{r} 11001000 = 200 \\ + 01100100 = 100 \\ \hline = 00101100 = 44 \end{array}$$

When represented with eight bits by a computer, 200 is stored as 11001000. Adding 100 results in an overflow, as the ninth bit of the result is discarded.

The meaning of operators such as “+” is therefore defined by means of *modular* arithmetic. However, the problem of reasoning about bit vectors extends beyond that of overflow and modular arithmetic. For efficiency reasons, programmers use bit-level operators to encode as much information as possible into the number of bits available.

As an example, consider the implementation of a propositional SAT solver. Recall the definition of a *literal* (Definition 1.11): a literal is a variable or its negation. Propositional SAT solvers that operate on formulas in CNF have to store a large number of such literals. We assume that we have numbered the variables that occur in the formula, and denote the variables by  $x_1, x_2, \dots$

The DIMACS standard for CNF uses signed numbers to encode a literal, e.g., the literal  $\neg x_3$  is represented as  $-3$ . The fact that we use signed numbers for the encoding avoids the use of one bit vector to store the sign. On the other hand, it reduces the possible number of variables to  $2^{31} - 1$  (the index 0 cannot be used any more), but this is still more than sufficient for any practical purpose.

In order to extract the index of a variable, we have to perform a case-split on the sign of the bit vector, for example as follows:

```
unsigned variable_index(int literal) {
    if(literal < 0)
        return -literal;
    else
        return literal;
}
```

The branch needed to implement the `if` statement in the program above slows down the execution of the program, as it is hard to predict for the branch prediction mechanisms of modern processors. Most SAT solvers therefore use a different encoding: the least significant bit of the bit vector is used to encode the sign of the literal, and the remaining bits encode the variable. The index of the variable can then be extracted by means of a bit-vector right-shift operation:

```
unsigned variable_index(unsigned literal) {
    return literal >> 1;
}
```

Similarly, the sign can be obtained by means of a bitwise AND operation:

```
bool literal_sign(unsigned literal) {
    return literal & 1;
}
```

The bitwise right-shift operation and the bitwise AND are implemented in most microprocessors, and both can be executed efficiently. Such bitwise operators also frequently occur in hardware design. Reasoning about such artifacts requires *bit-vector arithmetic*.

### 6.1.2 Notation

We use a simple variant of Church's  **$\lambda$ -Notation** in order to define vectors easily. A lambda expression for a bit vector with  $l$  bits has the form

$$\lambda i \in \{0, \dots, l - 1\}. f(i), \quad (6.2)$$

where  $f(i)$  is an expression that denotes the value of the  $i$ -th bit.

The use of the  $\lambda$ -operator to denote bit vectors is best explained by an example.

**Example 6.1.** Consider the following expressions.

- The expression

$$\lambda i \in \{0, \dots, l - 1\}. 0 \quad (6.3)$$

denotes the  $l$ -bit bit vector that consists only of zeros.

- A  $\lambda$ -expression is simply another way of defining a function *without* giving it a name. Thus, instead of defining a function  $z$  with


$$z(i) \doteq 0, \quad (6.4)$$

we can simply write  $\lambda i \in \{0, \dots, l - 1\}. 0$  for  $z$ .

- The expression

$$\lambda i \in \{0, \dots, 7\}. \begin{cases} 0 : i \text{ is even} \\ 1 : \text{otherwise} \end{cases} \quad (6.5)$$

denotes the bit vector 10101010.



**Fig. 6.1.** A bit vector  $b$  with  $l$  bits. The bit number  $i$  is denoted by  $b_i$

- The expression

$$\lambda i \in \{0, \dots, l-1\}. \neg x_i \quad (6.6)$$

denotes the bitwise negation of the vector  $x$ .

▀

We omit the domain of  $i$  from the lambda expression if the number of bits is clear from the context.

### 6.1.3 Semantics

We now give a formal definition of the meaning of a bit-vector arithmetic formula. We first clarify what a bit vector is.

**Definition 6.2 (bit vector).** A bit vector  $b$  is a vector of bits with a given length  $l$  (or dimension):

$$b : \{0, \dots, l-1\} \longrightarrow \{0, 1\}. \quad (6.7)$$

$bvec_l$

The set of all  $2^l$  bit vectors of length  $l$  is denoted by  $bvec_l$ . The  $i$ -th bit of the bit vector  $b$  is denoted by  $b_i$  (Fig. 6.1).

The meaning of a bit-vector formula obviously depends on the width of the bit-vector variables in it. This applies even if no arithmetic is used. As an example,

$$x \neq y \wedge x \neq z \wedge y \neq z \quad (6.8)$$

is unsatisfiable for bit vectors  $x$ ,  $y$ , and  $z$  that are one bit wide, but satisfiable for larger widths.

We sometimes use bit vectors that encode positive numbers only (unsigned bit vectors), and also bit vectors that encode both positive and negative numbers (signed bit vectors). Thus, each expression is associated with a **type**. The type of a bit-vector expression is

1. the width of the expression in bits, and
2. whether it is signed or unsigned.

We restrict the presentation to bit vectors that have a fixed, given length, as bit-vector arithmetic becomes undecidable as soon as arbitrary-width bit vectors are permitted. The width is known in most problems that arise in practice.

In order to clarify the type of an expression, we add indices in square brackets to the operator and operands in order to denote the bit-width (this is not to be confused with  $b_l$ , which denotes bit  $l$  of  $b$ ). As an example,  $a_{[32]} \cdot_{[32]} b_{[32]}$  denotes the multiplication of  $a$  and  $b$ . Both the result and the operands are 32 bits wide, and the remaining 32 bits of the result are discarded. The expression  $a_{[8]} \circ_{[24]} b_{[16]}$  denotes the concatenation of  $a$  and  $b$  and is in total 24 bits wide. In most cases, the width is clear from the context, and we therefore usually omit the subscript.

## Bitwise Operators

The meanings of bitwise operators can be defined through the bit vectors that they yield. The binary bitwise operators take two  $l$ -bit bit vectors as arguments and return an  $l$ -bit bit vector. As an example, the signature of the bitwise OR operator “ $|$ ” is

$$|_{[l]} : (bvec_l \times bvec_l) \longrightarrow bvec_l . \quad (6.9)$$

Using the  $\lambda$ -notation, the bitwise OR operator is defined as follows:

$$a | b \doteq \lambda i. (a_i \vee b_i) . \quad (6.10)$$

All the other bitwise operators are defined in a similar manner. In the following, we typically provide both the signature and the definition together.

## Arithmetic Operators

The meaning of a bit-vector formula with arithmetic operators depends on the *interpretation* of the bit vectors that it contains. There are many ways to encode numbers using bit vectors. The most commonly used encodings for integers are the **binary encoding** for unsigned integers and **two’s complement** for signed integers.

**Definition 6.3 (binary encoding).** *Let  $x$  denote a natural number, and  $b \in bvec_l$  a bit vector. We call  $b$  a binary encoding of  $x$  iff*

$$x = \langle b \rangle_U , \quad (6.11)$$

where  $\langle b \rangle$  is defined as follows:

$$\begin{aligned} \langle \cdot \rangle_U : bvec_l &\longrightarrow \{0, \dots, 2^l - 1\} , \\ \langle b \rangle_U &\doteq \sum_{i=0}^{l-1} b_i \cdot 2^i . \end{aligned} \quad (6.12)$$

The bit  $b_0$  is called the **least significant bit**, and the bit  $b_{l-1}$  is called the **most significant bit**.

$\langle \cdot \rangle_U$

Binary encoding can be used to represent non-negative integers only. One way of encoding negative numbers as well is to use one of the bits as a **sign bit**.

A naive way of using a sign bit is to simply negate the number if a designated bit is set, for example the most significant bit. As an example, 1001 could be interpreted as  $-1$  instead of 1. This encoding is hardly ever used in practice.<sup>1</sup> Instead, most microprocessor architectures implement the **two's complement** encoding.

**Definition 6.4 (two's complement).** *Let  $x$  denote a natural number, and  $b \in \text{bvec}_l$  a bit vector. We call  $b$  the two's complement of  $x$  iff*

$$x = \langle b \rangle_S, \quad (6.13)$$

where  $\langle \cdot \rangle_S$  is defined as follows:

$$\begin{aligned} \langle \cdot \rangle_S : \text{bvec}_l &\longrightarrow \{-2^{l-1}, \dots, 2^{l-1} - 1\}, \\ \langle b \rangle_S &:= -2^{l-1} \cdot b_{l-1} + \sum_{i=0}^{l-2} b_i \cdot 2^i. \end{aligned} \quad (6.14)$$

The bit with index  $l - 1$  is called the sign bit of  $b$ .

**Example 6.5.** Some encodings of integers in binary and two's complement are

$$\begin{aligned} \langle 11001000 \rangle_U &= 200, \\ \langle 11001000 \rangle_S &= -128 + 64 + 8 = -56, \\ \langle 01100100 \rangle_S &= 100. \end{aligned}$$

▀

Note that the meanings of the relational operators “>”, “<”, “≤”, “≥”, the multiplicative operators “·”, “/”, and the right-shift operator “>>” depend on whether a binary encoding or a two's complement encoding is used for the operands, which is why the encoding of the bit vectors is part of the type. We use the subscript  $U$  for a binary encoding (unsigned) and the subscript  $S$  for a two's complement encoding (signed). We may omit this subscript if the encoding is clear from the context, or if the meaning of the operator does not depend on the encoding (this is the case for most operators).

As suggested by the example at the beginning of this chapter, arithmetic on bit vectors has a wraparound effect: if the number of bits required to represent the result exceeds the number of bits available, the additional bits of the result are discarded, i.e., the result is truncated. This corresponds to a **modulo** operation, where the base is  $2^l$ . We write

$$x = y \pmod{b} \quad (6.15)$$

to denote that  $x$  and  $y$  are equal modulo  $b$ . The use of modulo arithmetic allows a straightforward definition of the interpretation of all arithmetic operators:

<sup>1</sup> The main reason for this is the fact that it makes the implementation of arithmetic operators such as addition more complicated, and that there are two encodings for 0, namely 0 and -0.

- Addition and subtraction:

$$a_{[l]} +_U b_{[l]} = c_{[l]} \iff \langle a \rangle_U + \langle b \rangle_U = \langle c \rangle_U \pmod{2^l}, \quad (6.16)$$

$$a_{[l]} -_U b_{[l]} = c_{[l]} \iff \langle a \rangle_U - \langle b \rangle_U = \langle c \rangle_U \pmod{2^l}, \quad (6.17)$$

$$a_{[l]} +_S b_{[l]} = c_{[l]} \iff \langle a \rangle_S + \langle b \rangle_S = \langle c \rangle_S \pmod{2^l}, \quad (6.18)$$

$$a_{[l]} -_S b_{[l]} = c_{[l]} \iff \langle a \rangle_S - \langle b \rangle_S = \langle c \rangle_S \pmod{2^l}. \quad (6.19)$$

Note that  $a +_U b = a +_S b$  and  $a -_U b = a -_S b$  (see Problem 6.7), and thus the  $U/S$  subscript can be omitted from the addition and subtraction operands. A semantics for mixed-type expressions is also easily defined, as shown in the following example:

$$a_{[l]U} +_U b_{[l]S} = c_{[l]U} \iff \langle a \rangle + \langle b \rangle_S = \langle c \rangle \pmod{2^l}. \quad (6.20)$$

- Unary minus:

$$-a_{[l]} = b_{[l]} \iff -\langle a \rangle_S = \langle b \rangle_S \pmod{2^l}. \quad (6.21)$$

- Relational operators:

$$a_{[l]U} < b_{[l]U} \iff \langle a \rangle_U < \langle b \rangle_U, \quad (6.22)$$

$$a_{[l]S} < b_{[l]S} \iff \langle a \rangle_S < \langle b \rangle_S, \quad (6.23)$$

$$a_{[l]U} < b_{[l]S} \iff \langle a \rangle_U < \langle b \rangle_S, \quad (6.24)$$

$$a_{[l]S} < b_{[l]U} \iff \langle a \rangle_S < \langle b \rangle_U. \quad (6.25)$$

The semantics for the other relational operators such as “ $\geq$ ” follows the same pattern. Note that ANSI-C compilers do not implement the relational operators on operands with mixed encodings the way they are formalized above (see Problem 6.6). Instead, the signed operand is converted to an unsigned operand, which does not preserve the meaning expected by many programmers.

- Multiplication and division:

$$a_{[l]} \cdot_U b_{[l]} = c_{[l]} \iff \langle a \rangle_U \cdot \langle b \rangle_U = \langle c \rangle_U \pmod{2^l}, \quad (6.26)$$

$$a_{[l]}/_U b_{[l]} = c_{[l]} \iff \langle a \rangle_U / \langle b \rangle_U = \langle c \rangle_U \pmod{2^l}, \quad (6.27)$$

$$a_{[l]} \cdot_S b_{[l]} = c_{[l]} \iff \langle a \rangle_S \cdot \langle b \rangle_S = \langle c \rangle_S \pmod{2^l}, \quad (6.28)$$

$$a_{[l]}/_S b_{[l]} = c_{[l]} \iff \langle a \rangle_S / \langle b \rangle_S = \langle c \rangle_S \pmod{2^l}. \quad (6.29)$$

- The extension operator: converting a bit vector to a bit vector with more bits is called **zero extension** in the case of an unsigned bit vector, and **sign extension** in the case of a signed bit vector. Let  $l \leq m$ . The value that is encoded does not change:

$$ext_{[m]U}(a_{[l]}) = b_{[m]U} \iff \langle a \rangle_U = \langle b \rangle_U, \quad (6.30)$$

$$ext_{[m]S}(a_{[l]}) = b_{[m]S} \iff \langle a \rangle_S = \langle b \rangle_S. \quad (6.31)$$

- **Shifting:** the left-shift operator “<<” takes two operands and shifts the first one to the left as many times as is given by the respective value of the second operand. The width of the left-hand-side operand is called the *width of the shift*, whereas the width of the right-hand-side operand is the *width of the shift-distance*. The vector is filled up with zeros from the right:

$$a_{[l]} \ll b_U = \lambda i \in \{0, \dots, l-1\}. \begin{cases} a_{i-\langle b \rangle} & : i \geq \langle b \rangle_U \\ 0 & : \text{otherwise} . \end{cases} \quad (6.32)$$

If  $b$  is a signed number, we replace  $\langle b \rangle_U$  by  $\langle b \rangle_S$ . See also Problem 6.5. The meaning of the right-shift “>>” operator depends on the encoding of the first operand: if it uses binary encoding (which, recall, is for unsigned bit vectors), zeros are inserted from the left. This is called a **logical right shift**:

$$a_{[l]_U} \gg b_U = \lambda i \in \{0, \dots, l-1\}. \begin{cases} a_{i+\langle b \rangle} & : i < l - \langle b \rangle \\ 0 & : \text{otherwise} . \end{cases} \quad (6.33)$$

If the first operand uses two’s complement encoding, the sign bit of  $a$  is replicated. This is also called an **arithmetic right shift**:

$$a_{[l]_S} \gg b_U = \lambda i \in \{0, \dots, l-1\}. \begin{cases} a_{i+\langle b \rangle} & : i < l - \langle b \rangle \\ a_{l-1} & : \text{otherwise} . \end{cases} \quad (6.34)$$

## 6.2 Deciding Bit-Vector Arithmetic with Flattening

### 6.2.1 Converting the Skeleton

The most commonly used decision procedure for bit-vector arithmetic is called *flattening*.<sup>2</sup> Algorithm 6.2.1 implements this technique. For a given bit-vector arithmetic formula  $\phi$ , the algorithm computes an equisatisfiable propositional formula  $\mathcal{B}$ , which is then passed to a SAT solver.

Let  $At(\phi)$  denote the set of atoms in  $\phi$ . As a first step, the algorithm replaces the atoms in  $\phi$  with new Boolean variables. We denote the variable that replaces an atom  $a \in At(\phi)$  by  $e(a)$ , and call this the **propositional encoder** of  $a$ . The resulting formula is denoted by  $e(\phi)$ . We call it the **propositional skeleton** of  $\phi$ . The propositional skeleton is the expression that is assigned to  $\mathcal{B}$  initially.

Let  $T(\phi)$  denote the set of terms in  $\phi$ . The algorithm then assigns a vector of new Boolean variables to each bit-vector term in  $T(\phi)$ . We use  $e(t)$  to denote this vector of variables for a given  $t \in T(\phi)$ , and  $e(t)_i$  to denote the variable for the bit with index  $i$  of the term  $t$ . The width of  $e(t)$  matches the width of the term  $t$ . Note that, so far, we have used  $e$  to denote three different, but related things: a propositional encoder of an atom, a propositional

<sup>2</sup> In colloquial terms, this technique is sometimes referred to as “*bit-blasting*”.

$\mathcal{B}$

$At(\phi)$

$e(\phi)$

$T(\phi)$

$e(t)$



formula resulting from replacing all atoms of a formula with their respective propositional encoders, and a propositional encoder of a term.

The algorithm then iterates over the terms and atoms of  $\phi$ , and computes a constraint for each of them. The constraint is returned by the function BV-CONSTRAINT, and is added as a conjunct to  $\mathcal{B}$ .

**Algorithm 6.2.1: BV-FLATTENING**

**Input:** A formula  $\phi$  in bit-vector arithmetic

**Output:** An equisatisfiable Boolean formula  $\mathcal{B}$

```

1. function BV-FLATTENING
2.    $\mathcal{B} := e(\phi)$ ; ▷ the propositional skeleton of  $\phi$ 
3.   for each  $t_{[l]} \in T(\phi)$  do
4.     for each  $i \in \{0, \dots, l-1\}$  do
5.       set  $e(t)_i$  to a new Boolean variable;
6.   for each  $a \in At(\phi)$  do
7.      $\mathcal{B} := \mathcal{B} \wedge \text{BV-CONSTRAINT}(e, a)$ ;
8.   for each  $t_{[l]} \in T(\phi)$  do
9.      $\mathcal{B} := \mathcal{B} \wedge \text{BV-CONSTRAINT}(e, t)$ ;
10.  return  $\mathcal{B}$ ;

```

The constraint that is needed for a particular atom  $a$  or term  $t$  depends on the atom or term, respectively. In the case of a bit vector or a Boolean variable, no constraint is needed, and BV-CONSTRAINT returns TRUE. If  $t$  is a bit-vector constant  $C_{[l]}$ , the following constraint is generated:

$$\bigwedge_{i=0}^{l-1} (C_i \iff e(t)_i). \quad (6.35)$$

Otherwise,  $t$  must contain a bit-vector operator. The constraint that is needed depends on this operator. The constraints for the bitwise operators are straightforward. As an example, consider bitwise OR, and let  $t = a_{[l]}b$ . The constraint returned by BV-CONSTRAINT is:

$$\bigwedge_{i=0}^{l-1} ((a_i \vee b_i) \iff e(t)_i). \quad (6.36)$$

The constraints for the other bitwise operators follow the same pattern.

### 6.2.2 Arithmetic Operators

The constraints for the arithmetic operators often follow implementations of these operators as a *circuit*. There is an abundance of literature on how to

build efficient circuits for various arithmetic operators. However, experiments with various alternative circuits have shown that the simplest ones usually burden the SAT solver the least. We begin by defining a one-bit adder, also called a **full adder**.

**Definition 6.6 (full adder).** A full adder is defined using the two functions *carry* and *sum*. Both of these functions take three input bits  $a$ ,  $b$ , and  $cin$  as arguments. The function *carry* calculates the carry-out bit of the adder, and the function *sum* calculates the sum bit:

$$sum(a, b, cin) \doteq (a \oplus b) \oplus cin, \quad (6.37)$$

$$carry(a, b, cin) \doteq (a \wedge b) \vee ((a \oplus b) \wedge cin). \quad (6.38)$$

We can extend this definition to adders for bit vectors of arbitrary length.

**Definition 6.7 (carry bits).** Let  $x$  and  $y$  denote two  $l$ -bit bit vectors and  $cin$  a single bit. The carry bits  $c_0$  to  $c_l$  are defined recursively as follows:

$$c_i \doteq \begin{cases} cin & : i = 0 \\ carry(x_{i-1}, y_{i-1}, c_{i-1}) & : otherwise. \end{cases} \quad (6.39)$$

**Definition 6.8 (adder).** An  $l$ -bit adder maps two  $l$ -bit bit vectors  $x$ ,  $y$  and a carry-in bit  $cin$  to their sum and a carry-out bit. Let  $c_i$  denote the  $i$ -th carry bit as in Definition 6.7. The function *add* is defined using the carry bits  $c_i$ :

$$add(x, y, cin) \doteq \langle result, cout \rangle, \quad (6.40)$$

$$result_i \doteq sum(x_i, y_i, c_i) \quad \text{for } i \in \{0, \dots, l-1\}, \quad (6.41)$$

$$cout \doteq c_n. \quad (6.42)$$

The circuit equivalent of this construction is called a *ripple carry adder*. One can easily implement the constraint for  $t = a + b$  using an adder with  $cin = 0$ :

$$\bigwedge_{i=0}^{l-1} (add(a, b, 0).result_i \iff e(t)_i). \quad (6.43)$$

One can prove by induction on  $l$  that (6.43) holds if and only if  $\langle a \rangle_U + \langle b \rangle_U = \langle e(t) \rangle_U \bmod 2^l$ , which shows that the constraint complies with the semantics.

Subtraction, where  $t = a - b$ , is implemented with the same circuit by using the following constraint (recall that  $\sim b$  is the bitwise negation of  $b$ ):

$$\bigwedge_{i=0}^{l-1} (add(a, \sim b, 1).result_i \iff e(t)_i). \quad (6.44)$$

This implementation makes use of the fact that  $\langle (\sim b) + 1 \rangle_S = -\langle b \rangle_S \bmod 2^l$  (see Problem 6.8).

## Relational Operators

The equality  $a =_{[l]} b$  is implemented using simply a conjunction:

$$\bigwedge_{i=0}^{l-1} a_i = b_i \iff e(t) . \quad (6.45)$$

The relation  $a < b$  is transformed into  $a - b < 0$ , and an adder is built for the subtraction, as described above. Thus,  $b$  is negated and the carry-in bit of the adder is set to TRUE. The result of the relation  $a < b$  depends on the encoding. In the case of unsigned operands,  $a < b$  holds if the carry-out bit *cout* of the adder is FALSE:

$$\langle a \rangle_U < \langle b \rangle_U \iff \neg \text{add}(a, \sim b, 1).cout . \quad (6.46)$$

In the case of signed operands,  $a < b$  holds if and only if  $(a_{l-1} = b_{l-1}) \neq \text{cout}$ :

$$\langle a \rangle_S < \langle b \rangle_S \iff (a_{l-1} \iff b_{l-1}) \oplus \text{add}(a, b, 1).cout . \quad (6.47)$$

Comparisons involving mixed encodings are implemented by extending both operands by one bit, followed by a signed comparison.

## Shifts

Recall that we call the width of the left-hand-side operand of a shift (the vector that is to be shifted) the *width of the shift*, whereas the width of the right-hand-side operand is the *width of the shift distance*.

We restrict the left and right shifts as follows: the width  $l$  of the shift must be a power of two, and the width of the shift distance  $n$  must be  $\log_2 l$ .

With this restriction, left and right shifts can be implemented by using the following construction, which is called the *barrel shifter*. The shifter is split into  $n$  stages. Stage  $s$  can shift the operand by  $2^s$  bits or leave it unaltered. The function  $ls$  is defined recursively for  $s \in \{-1, \dots, n-1\}$ :

$$ls(a_{[l]}, b_{[n]}_U, -1) \doteq a \quad (6.48)$$

$$ls(a_{[l]}, b_{[n]}_U, s) \doteq \lambda i \in \{0, \dots, l-1\}. \begin{cases} (ls(a, b, s-1))_{i-2^s} & : i \geq 2^s \wedge b_s \\ (ls(a, b, s-1))_i & : i \geq 2^s \wedge \neg b_s \\ 0 & : \text{otherwise} . \end{cases} \quad (6.49)$$

The barrel shifter construction needs only  $O(n \log n)$  logical operators, in contrast to the naive implementation, which requires  $O(n^2)$  operators.

## Multiplication and Division

Multipliers can be implemented following the most simplistic circuit design, which uses the *shift-and-add* idea. The function  $mul$  is defined recursively for  $s \in \{-1, \dots, n-1\}$ , where  $n$  denotes the width of the second operand:

$$mul(a, b, -1) \doteq b, \quad (6.50)$$

$$mul(a, b, s) \doteq mul(a, b, s-1) + (b_s?(a \ll s) : 0). \quad (6.51)$$

A division  $a/_{U}b$  is implemented by adding two constraints:

$$b \neq 0 \implies e(t) \cdot b + r = a. \quad (6.52)$$

$$b \neq 0 \implies r < b. \quad (6.53)$$

The variable  $r$  is a new bit vector of the same width as  $b$ , and contains the remainder. The signed-division and modulo operations are done in a similar way.

## 6.3 Incremental Bit Flattening

### 6.3.1 Some Operators Are Hard

For some operators, the size of the formula generated by BV-CONSTRAINT may be large. As an example, consider the formula for a single multiplier with  $n$  bits. The table in Fig. 6.2 shows the number of variables and the number of CNF clauses that are generated from the formula using Tseitin's encoding (see Sect. 1.3).

$n$	Number of variables	Number of clauses
8	313	1001
16	1265	4177
24	2857	9529
32	5089	17057
64	20417	68929

**Fig. 6.2.** The size of the constraint for an  $n$ -bit multiplier expression after Tseitin's transformation

In addition to the sheer size of these formulas, their symmetry and connectivity is a burden on the decision heuristic of state-of-the-art propositional SAT solvers. As a consequence, formulas with multipliers are often very hard to solve. Similar observations hold for other arithmetic operators such as division and modulo.

As an example, consider the following bit-vector formula:

$$a \cdot b = c \wedge b \cdot a \neq c \wedge x < y \wedge x > y \quad (6.54)$$

When this formula is encoded into CNF, a SAT instance with about 11 000 variables is generated for a width of 32 bits. This formula is obviously unsatisfiable. There are two reasons for this: the first two conjuncts are inconsistent, and independently, the last two conjuncts are inconsistent. The decision heuristics of most SAT solvers (see Chap. 2) are biased towards splitting first on variables that are used frequently, and thus favor decisions on  $a$ ,  $b$ , and  $c$ . Consequently, they attempt to show unsatisfiability of the formula on the hard part, which includes the two multipliers. The “easy” part of the formula, which contains only two relational operators, is ignored. Most propositional SAT solvers cannot solve this formula in a reasonable amount of time.

In many cases, it is therefore beneficial to build the flattened formula  $\mathcal{B}$  *incrementally*. Algorithm 6.3.1 is a realization of this idea: as before, we start with the propositional skeleton of  $\phi$ . We then add constraints for the “inexpensive” operators, and omit the constraints for the “expensive” operators. The bitwise operators are typically inexpensive, whereas arithmetic operators are expensive. The encodings with missing constraints can be considered an *abstraction* of  $\phi$ , and thus the algorithm is an instance of the abstraction–refinement procedure introduced in Sect. 3.4.

The current flattening  $\mathcal{B}$  is passed to a propositional SAT solver. If  $\mathcal{B}$  is unsatisfiable, so is the original formula  $\phi$ . Recall the formula (6.54): as soon as the constraints for the second half of the formula are added to  $\mathcal{B}$ , the encoding becomes unsatisfiable, and we may conclude that (6.54) is unsatisfiable without considering the multipliers.

On the other hand, if  $\mathcal{B}$  is satisfiable, one of two cases applies:

1. The original formula  $\phi$  is unsatisfiable, but one (or more) of the omitted constraints is needed to show this.
2. The original formula  $\phi$  is satisfiable.

In order to distinguish between these two cases, we can check whether the satisfying assignment produced by the SAT solver satisfies the constraints that we have omitted. As we might have removed variables, the satisfying assignment might have to be extended by setting the missing values to some constant, for example zero. If this assignment satisfies all constraints, the second case applies, and the algorithm terminates.

If this is not so, one or more of the terms for which the constraints were omitted is inconsistent with the assignment provided by the SAT solver. We denote this set of terms by  $I$ . The algorithm proceeds by selecting some of these terms, adding their constraints to  $\mathcal{B}$ , and reiterating. The algorithm terminates, as we strictly add more constraints with each iteration. In the worst case, all constraints from  $T(\phi)$  are added to the encoding.

**Algorithm 6.3.1: INCREMENTAL BV-FLATTENING****Input:** A formula  $\phi$  in bit-vector logic**Output:** “Satisfiable” if the formula is satisfiable, and “Unsatisfiable” otherwise

```

1. function INCREMENTAL-BV-FLATTENING( $\phi$ )
2.    $\mathcal{B} := e(\phi)$ ; ▷ propositional skeleton of  $\phi$ 
3.   for each  $t_{[l]} \in T(\phi)$  do
4.     for each  $i \in \{0, \dots, l-1\}$  do
5.       set  $e(t)_i$  to a new Boolean variable;
6.   while (TRUE) do
7.      $\alpha := \text{SAT-SOLVER}(\mathcal{B})$ ;
8.     if  $\alpha = \text{“Unsatisfiable”}$  then
9.       return “Unsatisfiable”;
10.    else
11.      Let  $I \subseteq T(\phi)$  be the set of terms that are inconsistent with the
        satisfying assignment;
12.      if  $I = \emptyset$  then
13.        return “Satisfiable”;
14.      else
15.        Select “easy”  $F' \subseteq I$ ;
16.        for each  $t_{[l]} \in F'$  do
17.           $\mathcal{B} := \mathcal{B} \wedge \text{BV-CONSTRAINT}(e, t)$ ;

```

**6.3.2 Enforcing Functional Consistency**

In many cases, omitting constraints for particular operators may result in a flattened formula that is too weak, and thus is satisfied by too many spurious models. On the other hand, the full constraint may burden the SAT solver too much. A compromise between the maximum strength of the full constraint and omitting the constraint altogether is to replace functions over bit-vectors by uninterpreted functions, and then reduce them to equalities while enforcing *functional consistency* only. The concept of functional consistency was presented in Chap. 3. This technique is particularly effective when one is checking the equivalence of two models.

For example, let  $a_1 op b_1$  and  $a_2 op b_2$  be two terms, where  $op$  is some binary operator (for simplicity, assume that these are the only terms in the input formula that use  $op$ ). First, replace  $op$  with a new uninterpreted-function symbol  $G$ . Second, apply Ackermann’s reduction in order to eliminate  $G$ : replace every occurrence of  $G(a_1, b_1)$  with a new variable  $g_1$ , and every occurrence of  $G(a_2, b_2)$  with a new variable  $g_2$ . Finally, add the functional-consistency constraint

$$a_1 = a_2 \wedge b_1 = b_2 \implies g_1 = g_2 . \quad (6.55)$$

The resulting formula does not contain constraints that correspond to the flattening of  $op$ . It is still necessary, however, to flatten the equalities resulting from the reduction.

## 6.4 Using Solvers for Linear Arithmetic

### 6.4.1 Motivation

The main disadvantage of flattening-based propositional encodings for formulas in bit-vector arithmetic is that all high-level structure present in the formula is lost. Another problem is that encoding an addition in propositional logic results in one XOR per bit. The XORs are chained together through the carry bit. It is known that such XOR chains can result in very hard SAT instances. As a result, there are many bit-vector formulas that cannot be decided by means of bit flattening and a SAT solver.

### 6.4.2 Integer Linear Arithmetic for Bit Vectors

We introduced decision procedures for linear arithmetic in Chap. 5. A restricted subset of bit-vector arithmetic can be translated into linear arithmetic over the integers to obtain a decision procedure that exploits the bit-vector structure (also known as the *word-level* structure) of the original decision problem.

**Definition 6.9 (linear bit-vector arithmetic).** *A term in bit-vector arithmetic that uses only constants on the right-hand side of binary bitwise, multiplication, and shift operators is called linear.*

We denote the linear atoms in a bit-vector formula  $\phi$  by  $A_L(\phi)$ , and the remaining atoms (the nonlinear atoms) by  $A_N(\phi)$ .

Let  $a$  be a linear atom. As preparation, we perform a number of transformations on the terms contained in  $a$ . We write  $\llbracket b \rrbracket$  for the transformation of any bit-vector arithmetic term  $b$ .

- Let  $b \gg d$  denote a bitwise right-shift term that is contained in  $a$ , where  $b$  is a term and  $d$  is a constant. It is replaced by  $\llbracket b \rrbracket / 2^{(d)}$ , i.e.,

$$\llbracket b \gg d \rrbracket \doteq \llbracket b \rrbracket / 2^{(d)}. \quad (6.56)$$

Bitwise left shifts are handled in a similar manner.

- The bitwise negation of a term  $b$  is replaced with  $-\llbracket b \rrbracket + 1$ :

$$\llbracket \sim b \rrbracket \doteq -\llbracket b \rrbracket + 1. \quad (6.57)$$

$$A_L(\phi)$$

$$A_N(\phi)$$

- A bitwise AND term  $b_{[l]} \& 1$ , where  $b$  is any term, is replaced by a new integer variable  $x$  subject to the following constraints over  $x$  and a second new integer variable  $\sigma$ :

$$0 \leq x \leq 1 \ \wedge \ [b] = 2\sigma + x \ \wedge \ 0 \leq \sigma < 2^{l-1} \quad (6.58)$$

A bitwise AND with other constants can be replaced using shifts. This can be optimized further by joining together groups of adjacent one-bits in the constant on the right-hand side.

- The bitwise OR is replaced with bitwise negation and bitwise AND.

We are now left with addition, subtraction, multiplication, and division.

As the next step, the division operators are removed from the constraints. As an example, the constraint  $a /_{[32]} 3 = b$  becomes  $a = b \cdot_{[34]} 3$ . Note that the bit-width of the multiplication has to be increased in order to take overflow into account. The operands  $a$  and  $b$  are sign-extended if signed, and zero-extended if unsigned. After this preparation, we can assume the following form of the atoms without loss of generality:

$$c_1 \cdot t_1 +_{[l]} c_2 \cdot t_2 \ op \ b, \quad (6.59)$$

where  $op$  is one of the relational operators as defined in Sect. 6.1,  $c_1$ ,  $c_2$ , and  $b$  are constants, and  $t_1$  and  $t_2$  are bit-vector identifiers with  $l$  bits. Sums with more than two addends can be handled in a similar way.

As we can handle additions efficiently, all scalar multiplications  $c \cdot_{[l]} a$  with a small constant  $c$  are replaced by  $c$  additions. For example,  $3 \cdot a$  becomes  $a + a + a$ . For large coefficients, this is inefficient, and a different encoding is used: let  $\sigma$  be a new variable. The scalar multiplication is replaced by  $c \cdot a - 2^l \cdot \sigma$  together with the following constraints:

$$c \cdot a - 2^l \cdot \sigma \leq 2^l - 1 \ \wedge \ \sigma \leq c - 1. \quad (6.60)$$

### Case-Splitting for Overflow

After this transformation, we are left with bit-vector additions of the following form:

$$t_1 +_{[l]} t_2 \ op \ b. \quad (6.61)$$

If the constraints are passed in this form to a decision procedure for integer linear arithmetic, for example the Omega test, the potential overflow in the  $l$ -bit bit-vector addition is disregarded. Given that  $t_1$  and  $t_2$  are  $l$ -bit unsigned vectors, we have  $t_1 \in \{0, \dots, 2^l - 1\}$  and  $t_2 \in \{0, \dots, 2^l - 1\}$ , and, thus,  $t_1 + t_2 \in \{0, \dots, 2^{l+1} - 2\}$ . We use a case-split to adjust the value of the sum in the case of an overflow and transform (6.61) into

$$((t_1 + t_2 \leq 2^l - 1) ? t_1 + t_2 : (t_1 + t_2 - 2^l)) \ op \ b. \quad (6.62)$$



The Omega test does not itself handle the resulting case-splits, but the case-splits can be lifted up to the propositional level by introducing an additional propositional variable  $p$ , and adding the following constraints:

$$p \iff (t_1 + t_2 \leq 2^l - 1), \quad (6.63)$$

$$p \implies (t_1 + t_2) \text{ op } b, \quad (6.64)$$

$$\neg p \implies (t_1 + t_2 - 2^l) \text{ op } b. \quad (6.65)$$

Thus, the price paid for the bit-vector semantics is two additional integer constraints for each bit-vector addition in the original problem. In practice, this technique is known to perform well on problems in which most constraints are conjoined, but deteriorates on problems with a complex Boolean structure. The performance also suffers when many bitwise operators are used.

**Example 6.10.** Consider the following formula:

$$x_{[8]} +_{[8]} 100 \leq 10_{[8]}. \quad (6.66)$$

This formula is already in the form given by (6.61). We only need to add the case-split:

$$0 \leq x \leq 255, \quad (6.67)$$

$$p \iff (x + 100 \leq 255), \quad (6.68)$$

$$p \implies (x + 100) \leq 10, \quad (6.69)$$

$$\neg p \implies (x + 100 - 256) \leq 10. \quad (6.70)$$


The conjunction of (6.67)–(6.70) has satisfying assignments, one of which is  $\{p \mapsto \text{FALSE}, x \mapsto 160\}$ . This is also a satisfying assignment for (6.66).  $\blacksquare$

## 6.5 Fixed-Point Arithmetic

### 6.5.1 Semantics

Many applications, for example in scientific computing, require arithmetic on numbers with a fractional part. High-end microprocessors offer support for **floating-point arithmetic** for this purpose. However, fully featured floating-point arithmetic is too heavyweight for many applications, such as control software embedded in vehicles, and computer graphics. In these domains, **fixed-point arithmetic** is a reasonable compromise between accuracy and complexity. Fixed-point arithmetic is also commonly supported by database systems, for example to represent amounts of currency.

In fixed-point arithmetic, the representation of the number is partitioned into two parts, the *integer part* (also called the *magnitude*) and the *fractional part* (Fig. 6.3). The number of digits in the fractional part is fixed – hence the



**Fig. 6.3.** A fixed-point bit vector  $b$  with a total of  $j + k = l$  bits. The dot is called the radix point. The  $j$  bits before the dot represent the magnitude (the integer part), whereas the  $k$  bits after the dot represent the fractional part

name “fixed point arithmetic”. The number 1.980, for example, is a fixed-point number with a three-digit fractional part.

The same principle can be applied to binary arithmetic, as captured by the following definition. Recall the definition of  $\langle \cdot \rangle_S$  (two’s complement) from Sect. 6.1.3.

**Definition 6.11.** *Given two bit vectors  $M$  and  $F$  with  $m$  and  $f$  bits, respectively, we define the rational number that is represented by  $M.F$  as follows and denote it by  $\langle M.F \rangle$ :*

$$\begin{aligned} \langle \cdot \rangle &: \{0, 1\}^{m+f} \longrightarrow \mathbb{Q}, \\ \langle M.F \rangle &:= \frac{\langle M \circ F \rangle_S}{2^f}. \end{aligned}$$

**Example 6.12.** Some encodings of rational numbers as fixed-point numbers with base 2 are:

$$\begin{aligned} \langle 0.10 \rangle &= 0.5, \\ \langle 0.01 \rangle &= 0.25, \\ \langle 01.1 \rangle &= 1.5, \\ \langle 1111111.1 \rangle &= -0.5. \end{aligned}$$

Some rational numbers are not precisely representable using fixed-point arithmetic in base 2: they can only be approximated. As an example, for  $m = f = 4$ , the two numbers that are closest to  $1/3$  are

$$\begin{aligned} \langle 0000.0101 \rangle &= 0.3125, \\ \langle 0000.0110 \rangle &= 0.375. \end{aligned}$$

▀

Definition 6.11 gives us the semantics of fixed-point arithmetic. For example, the meaning of addition on bit vectors that encode fixed-point numbers can be defined as follows:

$$\begin{aligned} a_M.a_F + b_M.b_F = c_M.c_F &\iff \\ \langle a_M.a_F \rangle \cdot 2^f + \langle b_M.b_F \rangle \cdot 2^f &= \langle c_M.c_F \rangle \cdot 2^f \pmod{2^{m+f}}. \end{aligned}$$

There are variants of fixed-point arithmetic that implement **saturation** instead of overflow semantics, that is, instead of wrapping around, the result remains at the highest or lowest number that can be represented with the given precision. Both the semantics and the flattening procedure are straightforward for this case.

### 6.5.2 Flattening

Fixed-point arithmetic can be flattened just as well as arithmetic using binary encoding or two's complement. We assume that the numbers on the left- and right-hand sides of a binary operator have the same numbers of bits, before and after the radix point. If this is not so, missing bits after the radix point can be added by padding the fractional part with zeros from the right. Missing bits before the radix point can be added from the left using sign-extension.

The operators are encoded as follows:

- The bitwise operators are encoded exactly as in the case of binary numbers. Addition, subtraction, and the relational operators can also be encoded as in the case of binary numbers.
- Multiplication requires an alignment. The result of a multiplication of two numbers with  $f_1$  and  $f_2$  bits in the fractional part, respectively, is a number with  $f_1 + f_2$  bits in the fractional part. Note that, most commonly, fewer bits are needed, and thus, the extra bits of the result have to be rounded off using a separate **rounding** step.

**Example 6.13.** Addition and subtraction are straight-forward, but note the need for sign-extension in the second sum:

$$\begin{aligned}\langle 00.1 \rangle + \langle 00.1 \rangle &= \langle 01.0 \rangle \\ \langle 000.0 \rangle + \langle 1.0 \rangle &= \langle 111.0 \rangle\end{aligned}$$

The following examples illustrate multiplication without any subsequent rounding:

$$\begin{aligned}\langle 0.1 \rangle \cdot \langle 1.1 \rangle &= \langle 0.11 \rangle \\ \langle 1.10 \rangle \cdot \langle 1.1 \rangle &= \langle 10.010 \rangle\end{aligned}$$

If needed, rounding towards zero, towards the next even number, or towards  $+/-\infty$  can be applied in order to reduce the size of the fractional part; see Problem 6.9. ▀

There are many other encodings of numbers, which we do not cover here, e.g., binary-coded decimals (BCDs), or fixed-point formats with sign bit.

## 6.6 Problems

### 6.6.1 Semantics

**Problem 6.1 (operators that depend on the encoding).** Provide an example (with values of operands) that illustrates that the semantics depend on the encoding (signed vs. unsigned) for each of the following three operators:  $>$ ,  $\otimes$ , and  $\gg$ .

**Problem 6.2 ( $\lambda$ -notation).** Define the meaning of  $a_l \circ b_l$  using the  $\lambda$ -notation.

**Problem 6.3 (negation).** What is  $-10000000_S$  if the operand of the unary minus is a bit-vector constant?

**Problem 6.4 ( $\lambda$ -notation).** Define the meaning of  $a_{[l]U} \gg_{[l]U} b_{[m]S}$  and  $a_{[l]S} \gg_{[l]S} b_{[m]S}$  using modular arithmetic. Prove these definitions to be equivalent to the definition given in Sect. 6.1.3.

**Problem 6.5 (shifts in hardware).** What semantics of the left-shift does the processor in your computer implement? You can use a program to test this, or refer to the specification of the CPU. Formalize the semantics.

**Problem 6.6 (relations in hardware).** What semantics of the  $<$  operator does the processor in your computer implement if a signed integer is compared with an unsigned integer? Try this for the ANSI-C types `int`, `unsigned`, `char`, and `unsigned char`. Formalize the semantics, and specify the vendor and model of the CPU.

**Problem 6.7 (two's complement).** Prove

$$a_{[l] + U} b_{[l]} = a_{[l]} +_S b_{[l]}. \quad (6.71)$$

### 6.6.2 Bit-Level Encodings of Bit-Vector Arithmetic

**Problem 6.8 (negation).** Prove  $\langle (\sim b) + 1 \rangle_S = -\langle b \rangle_S \pmod{2^l}$ .

**Problem 6.9 (relational operators).** Prove the correctness of the flattening for “ $<$ ” as given in Sect. 6.2, for

- (a) unsigned operands,
- (b) signed operands,
- (c) an unsigned and a signed operand.

**Problem 6.10 (rounding for fixed-point arithmetic).** Formally specify the operator for rounding a fixed-point number with a fractional part of size  $f_1$  to a fractional part of size  $f_2 < f_1$  for the following cases:

- (a) rounding to zero,
- (b) rounding to  $-\infty$ , and
- (c) rounding to the nearest even number.

**Problem 6.11 (flattening fixed-point arithmetic).** Provide a flattening for the three rounding operators above.

### 6.6.3 Using Solvers for Linear Arithmetic

**Problem 6.12 (bitwise AND).** Give a translation of

$$x_{[32]U} = y_{[32]U} \& 0\text{xffff}0000 \quad (6.72)$$

into disjunctive integer linear arithmetic that is more efficient than that suggested by (6.58).

**Problem 6.13 (addition without splitting).** Can you propose a different translation for addition that does not use case-splitting but uses a new integer variable instead?

## 6.7 Bibliographic Notes

Bit-vector arithmetic was identified as an important logic for verification and equivalence checking in the hardware industry in [183]. The notation we use to annotate the type of the bit-vector expressions is taken from [32].

Early decision procedures for bit-vector arithmetic can be found in tools such as SVC [10] and ICS [74]. ICS used BDDs in order to decide properties of arithmetic operators, whereas SVC was based on a computation of a canonizer and a solver [12]. SVC has been superseded by CVC, and then CVC-Lite [9] and STP, both of which use a propositional SAT solver to decide the satisfiability of a circuit-based flattening of a bit-vector formula. ICS was superseded by YICES, which also uses flattening and a SAT solver.

Bounded model checking (BMC) is a common source of bit-vector arithmetic decision problems [23]. BMC was designed originally for synchronous models, as frequently found in the hardware domain, for example. BMC has been adopted in other domains that result in bit-vector formulas, for example software programs given in ANSI-C [48].

Translations to integer linear arithmetic have been used for bit-vector decision problems found in the hardware verification domain. Brinkmann and Drechsler [32] translated a fragment of bit-vector arithmetic into ILP and used the Omega test as a decision procedure for the ILP problem. However, the work in [32] was aimed only at the data-paths, and thus did not allow a Boolean part within the original formula. This was mended by Parthasarathy et al. [147] using an incremental encoding similar to the one described in Chap. 11.

COGENT [49] decides the validity of ANSI-C expressions. ANSI-C expressions are drawn from a fragment of bit-vector arithmetic, extended with pointer logic (see Chap. 8). COGENT and related procedures have many applications besides checking verification conditions. As an example, see [25, 26] for an application of Cogent to database testing. In addition to deciding the validity of ANSI-C expressions, C32SAT [33], developed by Brummayer and

Biere, is also able to determine if an expression always has a well-defined meaning according to the ANSI-C standard.

Current state-of-the-art decision procedures for bit-vector arithmetic apply heavy preprocessing to the formula, but ultimately rely on flattening a formula to propositional SAT [34, 123]. The preprocessing is especially beneficial if the formula also contains large arrays, for example for modeling memories [78, 122]. The tool SPEAR [7], by Babic and Hu, which is based on bit-blasting and a fast SAT solver with numerous optimization parameters that were tuned *automatically*, won the 2007 competition in the bit-vector category.

## 6.8 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$c?a : b$	Case-split on condition $c$	149
$\lambda$	Lambda expressions	151
$bvec_l$	Set of bit vectors with $l$ bits	152
$\langle \cdot \rangle_U$	Number encoded by binary encoding	153
$\langle \cdot \rangle_S$	Number encoded by two's complement	154
$A(\phi)$	Set of atoms in $\phi$	156
$T(\phi)$	Set of terms in $\phi$	156
$c_i$	Carry bit $i$	158

## Arrays

### 7.1 Introduction

The array is an important type of data structure that is used in most software programs, as well as in other domains, such as in modeling memories and caches in hardware design. This chapter introduces decision procedures for array logic. We focus on methods that perform a reduction from array logic to theories that we have already introduced.

Array logic permits expressions over arrays, which are formalized as maps from an *index type* to an *element type*. We denote the index type by  $T_I$ , and the element type by  $T_E$ . The type of the arrays themselves is denoted by  $T_A$ , which is a shorthand for  $T_I \rightarrow T_E$ , i.e., the set of functions that map an element of  $T_I$  to an element of  $T_E$ .

Let  $a \in T_A$  denote an array. There are two basic operations on arrays:

1. *Reading* an element with index  $i \in T_I$  from  $a$ . The value of the element that has index  $i$  is denoted by  $a[i]$ .
2. *Writing* an element with index  $i \in T_I$ . Let  $e \in T_E$  denote the value to be written. The array  $a$  where element  $i$  has been replaced by  $e$  is denoted by  $a\{i \leftarrow e\}$ .

We call the theories used to reason about the indices and the elements the *index theory* and the *element theory*, respectively.

The index logic should permit existential and universal quantification, in order to model properties such as “there exists an array element that is zero” or “all elements of the array are greater than zero”. An example of a suitable index logic is Presburger arithmetic, i.e., linear arithmetic over integers (Chap. 5) with quantification (Chap. 9). We can obtain multi-dimensional arrays by recursively defining  $T_A(n)$  for  $n$ -dimensional arrays. For  $n \geq 2$ , we simply add  $T_A(n - 1)$  to the element type of  $T_A(n)$ .

We start with a very general definition of array logic. Validity for this logic is not decidable, however, and we therefore add restrictions later on.

**Definition 7.1 (array logic).** *The syntax of a formula in array logic is defined by extending the syntactic rules for the index logic and the element logic. Let  $atom_I$  and  $atom_E$  denote an atom in the index logic and element logic, respectively, and let  $term_I$  and  $term_E$  denote a term in the index and element logic, respectively:*

$$\begin{aligned} atom &: atom_I \mid atom_E \mid \neg atom \mid atom \wedge atom \mid \\ &\quad \forall \text{array-identifier. } atom \\ term_A &: \text{array-identifier} \mid term_A\{term_I \leftarrow term_E\} \\ term_E &: term_A[term_I] \end{aligned}$$

Observe that equality between arrays is not permitted by the grammar. Equality between arrays  $a_1$  and  $a_2$  can be written as  $\forall i. a_1[i] = a_2[i]$ , assuming equality is permitted by the element theory.

The main axiom used to define the meanings of the two new operators above is the **read-over-write axiom**: after the value  $e$  has been written into array  $a$  at index  $i$ , the value of this array at index  $i$  is  $e$ . The value at any index  $j \neq i$  matches that in the array before the write operation at index  $j$ :

$$\forall a \in T_A. \forall e \in T_E. \forall i, j \in T_I. a\{i \leftarrow e\}[j] = \begin{cases} e & : i = j \\ a[j] & : \text{otherwise} \end{cases}. \quad (7.1)$$

As mentioned above, the problem of deciding the validity of an arbitrary formula in array logic is undecidable, even if the combination of the index logic and the element logic is decidable (see Problem 7.2). The following example illustrates the use of array logic for verifying an invariant.

**Example 7.2.** To illustrate the use of array logic in program verification, consider the pseudocode fragment in Fig. 7.1. The main step of the correctness argument is to show that the invariant in line 7 is maintained by the assignment in line 6. A common way to do so is to generate **verification conditions**, e.g., using Hoare's axiom system. We obtain the following verification condition for the claim:

$$\begin{aligned} &(\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0) \\ &\wedge a' = a\{i \leftarrow 0\} \\ \implies &(\forall x \in \mathbb{N}_0. x \leq i \implies a'[x] = 0). \end{aligned} \quad (7.2)$$

Proving validity of this formula shows that the loop invariant is maintained. This claim can be proven manually by means of the axiom in (7.1). We aim at an automatic procedure to decide the validity of expressions such as the one above. ▀

## 7.2 Arrays as Uninterpreted Functions

Consider the fragment of array logic which does not permit quantification over arrays, i.e., arrays are ground terms. A trivial way to reduce such formulas to a



```

1  a: array 0..99 of integer;
2  i: integer;
3
4  for i:=0 to 99 do
5      /*  $\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0$  */
6      a[i]:=0;
7      /*  $\forall x \in \mathbb{N}_0. x \leq i \implies a[x] = 0$  */
8  done;
9  /*  $\forall x \in \mathbb{N}_0. x \leq 99 \implies a[x] = 0$  */

```

**Fig. 7.1.** Pseudocode fragment that initializes an array of size 100 with zeros, annotated with the invariants that are maintained

### Aside: Array Bounds Checking in Programs

While array logic uses arrays of unbounded size, array data structures in programs are of bounded size. If an index variable exceeds the size of an array in a program, the value returned may be undefined or a crash might occur. This situation is called an **array bounds violation**. In the case of a write operation, other data might be overwritten, which is often exploitable to gain control over a computer system from a remote location over a network. Checking that a program never violates any of its array bounds is therefore highly desirable.

Note, however, that the issue of array bounds checking in programs does not require array logic; the question of whether an array index is within the bounds of a finite-size array requires one only to keep track of the *size* of the array, not of its contents.

As an example, consider the following program fragment, which is meant to move the elements of an array:

```

int a[N];

for(int i=0; i<N; i++)
    a[i]=a[i+1];

```

Despite of the fact that the program contains an array, the verification condition for the array-bounds property does not require array logic:

$$i < N \implies (i < N \wedge i + 1 < N). \quad (7.3)$$

combination of other theories is to treat the arrays as *uninterpreted functions* (see Chap. 3). The index operator is replaced by a function application, where the array is the (uninterpreted) function, and the index is the only function argument.

**Example 7.3.** Consider the following array logic formula, where  $a$  is an array with element type `char`:

$$(i = j \wedge a[j] = 'z') \implies a[i] = 'z' . \quad (7.4)$$

The character constant 'z' can be read as an integer number. Let  $F_a$  denote the uninterpreted function introduced for the array  $a$ :

$$(i = j \wedge F_a(j) = 'z') \implies F_a(i) = 'z' . \quad (7.5)$$

By applying Bryant's reduction (Chap. 3), we obtain

$$(i = j \wedge F_1^* = 'z') \implies F_2^* = 'z' , \quad (7.6)$$

where

$$F_1^* = f_1 \quad \text{and} \quad F_2^* = \begin{cases} f_1 : i = j \\ f_2 : \text{otherwise} . \end{cases} \quad (7.7)$$

The formula can then be shown to be valid by means of a decision procedure for equality (Chap. 4).  $\blacksquare$

Array updates can be handled by replacing each expression of the form  $a\{i \leftarrow e\}$  by a fresh variable  $a'$  of array type, and by adding two constraints that correspond directly to the two cases of the read-over-write axiom:

1.  $a'[i] = e$  for the value that is written,
2.  $\forall j \neq i. a'[j] = a[j]$  for the values that are unchanged.

This is called the **write rule**, and is an equivalence-preserving transformation on array logic formulas.

**Example 7.4.** The formula

$$a\{i \leftarrow e\}[i] \geq e \quad (7.8)$$

is transformed by introducing a new array identifier  $a'$  to replace  $a\{i \leftarrow e\}$ . Additionally, we add the assumption that  $a'[i] = e$ , and obtain

$$a'[i] = e \implies a'[i] \geq e , \quad (7.9)$$

which shows the validity of (7.8). The second part of the read-over-write axiom is needed to show the validity of a formula such as

$$a[0] = 10 \implies a\{1 \leftarrow 20\}[0] = 10 . \quad (7.10)$$

As before, the formula is transformed by replacing  $a\{1 \leftarrow 20\}$  with a new identifier  $a'$  and adding the two constraints described above:

$$(a[0] = 10 \wedge a'[1] = 20 \wedge (\forall j \neq 1. a'[j] = a[j])) \implies a'[0] = 10 . \quad (7.11)$$

Again as before, we transform this formula by replacing  $a$  and  $a'$  with uninterpreted-function symbols  $F_a$  and  $F_{a'}$ :

$$(F_a(0) = 10 \wedge F_{a'}(1) = 20 \wedge (\forall j \neq 1. F_{a'}(j) = F_a(j))) \implies F_{a'}(0) = 10 .$$

This simple example shows that array logic can be reduced to combinations of the index logic and uninterpreted functions, provided that the index logic offers quantifiers. The combination of Presburger arithmetic and uninterpreted functions is in general undecidable, however, and, thus, we need to restrict the set of formulas we consider. This is also the basic idea of the reduction algorithm in the following section.  $\blacksquare$

## 7.3 A Reduction Algorithm for Array Logic

### 7.3.1 Array Properties

We define here a restricted class of array logic formulas in order to obtain decidability. We consider formulas that are Boolean combinations of **array properties**.

**Definition 7.5 (array property).** *An array logic formula is called an array property if and only if it is of the form*

$$\forall i_1, \dots, i_k \in T_I. \phi_I(i_1, \dots, i_k) \implies \phi_V(i_1, \dots, i_k), \quad (7.12)$$

and satisfies the following conditions:

1. The predicate  $\phi_I$ , called the index guard, must follow the grammar

$$\begin{aligned} \text{iguard} &: \text{iguard} \wedge \text{iguard} \mid \text{iguard} \vee \text{iguard} \mid \text{iterm} \leq \text{iterm} \mid \text{iterm} = \text{iterm} \\ \text{iterm} &: i_1 \mid \dots \mid i_k \mid \text{term} \\ \text{term} &: \text{integer-constant} \mid \text{integer-constant} \cdot \text{index-identifier} \mid \text{term} + \text{term} \end{aligned}$$

The “index-identifier” used in “term” must not be one of  $i_1, \dots, i_k$ .

2. The index variables  $i_1, \dots, i_k$  can only be used in array read expressions of the form  $a[i_j]$ .

The predicate  $\phi_V$  is called the value constraint.

**Example 7.6.** The **extensionality rule** defines the equality of two arrays  $a_1$  as elementwise equality. Extensionality is an array property:

$$\forall i. a_1[i] = a_2[i]. \quad (7.13)$$

Note that the index guard is simply TRUE in this case.

Recall the array logic formula (7.2). The first and the third conjunct are obviously array properties, but recall the second conjunct,

$$a' = a\{i \leftarrow 0\}. \quad (7.14)$$

Is this an array property as well? As illustrated in Example 7.3, an array update expression can be replaced by adding two constraints. In our example, the first constraint is  $a'[i] = 0$ , which is obviously an array property. The second constraint is

$$\forall j \neq i. a'[j] = a[j], \quad (7.15)$$

and does not comply with the syntax constraints for index guards as given in Definition 7.5. However, it can be rewritten as

$$\forall j. (j \leq i - 1 \vee i + 1 \leq j) \implies a'[j] = a[j] \quad (7.16)$$

to match the syntactic constraints. ▀

### 7.3.2 A Reduction Algorithm

We now describe an algorithm that accepts a formula from the array property fragment of array logic and reduces it to an equisatisfiable formula that uses the element and index theories. The resulting formula can be further reduced to propositional logic using the methods described so far.

We assume that the following operators are defined for the index and element theories, and that we have a decision procedure for the combined theory:

- For the index type, we assume that linear arithmetic over indices is permitted.
- For the element type, we assume only that equality between two elements is permitted.

Algorithm 7.3.1 takes an array logic formula from the array property fragment as input. Note that the transformation of array properties to NNF may turn a universal quantification over the indices into an existential quantification. The formula does not contain explicit quantifier alternations, owing to the syntactic restrictions.

As a first step, the algorithm applies the write rule (see Sect. 7.2) to remove all array update operators. The resulting formulas contain quantification over indices, array reads, and subformulas from the element and index theories.

As the formula is in NNF, an existential quantification can be replaced by a new variable (which is implicitly existentially quantified). The universal quantifiers  $\forall i \in T_I. P(i)$  are replaced by the conjunction  $\bigwedge_{i \in \mathcal{I}(\phi)} P(i)$ , where the set  $\mathcal{I}(\phi)$  denotes the index expressions that  $i$  might possibly be equal to in the formula  $\phi$ . This set contains the following elements:

1. All expressions used as an array index in  $\phi$  that are not quantified variables.
2. All expressions used inside index guards in  $\phi$  that are not quantified variables.
3. If  $\phi$  contains none of the above,  $\mathcal{I}(\phi)$  is  $\{0\}$  in order to obtain a nonempty set of index expressions.

$\mathcal{I}(\phi)$

Finally, the array read operators are replaced by uninterpreted functions, as described in Sect. 7.2.

**Algorithm 7.3.1:** ARRAY-REDUCTION

**Input:** An array property formula  $\phi_A$  in NNF

**Output:** A formula  $\phi_{UF}$  in the index and element theories with uninterpreted functions

1. Apply the write rule to remove all array updates from  $\phi_A$ .
2. Replace all existential quantifications of the form  $\exists i \in T_I. P(i)$  by  $P(j)$ , where  $j$  is a fresh variable.
3. Replace all universal quantifications of the form  $\forall i \in T_I. P(i)$  by

$$\bigwedge_{i \in \mathcal{I}(\phi)} P(i).$$

4. Replace the array read operators by uninterpreted functions and obtain  $\phi_{UF}$ ;
5. **return**  $\phi_{UF}$ ;

**Example 7.7.** In order to illustrate Algorithm 7.3.1, we continue the introductory example by proving the validity of (7.2):

$$\begin{aligned} & (\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0) \\ & \wedge a' = a\{i \leftarrow 0\} \\ \implies & (\forall x \in \mathbb{N}_0. x \leq i \implies a'[x] = 0). \end{aligned}$$

That is, we are checking satisfiability of

$$\begin{aligned} & (\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0) \\ & \wedge a' = a\{i \leftarrow 0\} \\ & \wedge (\exists x \in \mathbb{N}_0. x \leq i \wedge a'[x] \neq 0). \end{aligned} \tag{7.17}$$

By applying the write rule, we obtain

$$\begin{aligned} & (\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0) \\ & \wedge a'[i] = 0 \wedge \forall j \neq i. a'[j] = a[j] \\ & \wedge (\exists x \in \mathbb{N}_0. x \leq i \wedge a'[x] \neq 0). \end{aligned} \tag{7.18}$$

In the second step of Algorithm 7.3.1, we instantiate the existential quantifier with a new variable  $z \in \mathbb{N}_0$ :

$$\begin{aligned} & (\forall x \in \mathbb{N}_0. x < i \implies a[x] = 0) \\ & \wedge a'[i] = 0 \wedge \forall j \neq i. a'[j] = a[j] \\ & \wedge z \leq i \wedge a'[z] \neq 0). \end{aligned} \tag{7.19}$$

The set  $\mathcal{I}$  for our example is  $\{i, z\}$ . We therefore replace the two universal quantifications as follows:

$$\begin{aligned} & (i < i \implies a[i] = 0) \wedge (z < i \implies a[z] = 0) \\ & \wedge a'[i] = 0 \wedge (i \neq i \implies a'[i] = a[i]) \wedge (z \neq i \implies a'[z] = a[z]) \\ & \wedge z \leq i \wedge a'[z] \neq 0). \end{aligned} \quad (7.20)$$

Let us remove the trivially satisfied conjuncts to obtain

$$\begin{aligned} & (z < i \implies a[z] = 0) \\ & \wedge a'[i] = 0 \wedge (z \neq i \implies a'[z] = a[z]) \\ & \wedge z \leq i \wedge a'[z] \neq 0). \end{aligned} \quad (7.21)$$

We now replace the two arrays  $a$  and  $a'$  by uninterpreted functions  $F_a$  and  $F_{a'}$  and obtain

$$\begin{aligned} & (z < i \implies F_a(z) = 0) \\ & \wedge F_{a'}(i) = 0 \wedge (z \neq i \implies F_{a'}(z) = F_a(z)) \\ & \wedge z \leq i \wedge F_{a'}(z) \neq 0). \end{aligned} \quad (7.22)$$

By distinguishing the three cases  $z < i$ ,  $z = i$ , and  $z > i$ , it is easy to see that this formula is unsatisfiable.  $\blacksquare$

## 7.4 Problems

**Problem 7.1 (manual proofs for array logic).** Show the validity of (7.2) using the read-over-write axiom.

**Problem 7.2 (undecidability of array logic).** Show that the satisfiability of an array logic formula is undecidable by using a reduction of a two-counter machine to an array logic formula: given a two-counter machine  $M$ , generate an array logic formula  $\varphi$  that is valid if  $M$  terminates.

**Problem 7.3 (quantifiers and NNF).** The transformation steps 3 and 4 of Algorithm 7.3.1 rely on the fact that the formula is in NNF. Provide one example for each of these steps that shows that the step is unsound if the formula is not in NNF.

## 7.5 Bibliographic Notes

The read-over-write axiom (7.1) is due to John McCarthy, who used it to show the correctness of a compiler for arithmetic expressions [124]. The reads and writes correspond to loads and stores in a computer memory. Hoare and Wirth introduced the notation  $(a, i : e)$  for  $a\{i \leftarrow e\}$ , and used it to define

the meaning of assignments to array elements in the PASCAL programming language [93].

Automatic decision procedures for arrays have been found in automatic theorem provers since the very beginning. In the context of program verification, array logic is often combined with application-specific predicates, for example to specify properties such as “the array is sorted” or to specify ranges of indices [164]. Greg Nelson’s theorem prover Simplify [65] has McCarthy’s read-over-write axiom and appropriate instantiation heuristics built in.

The reduction of array logic to fragments of Presburger arithmetic with uninterpreted functions is commonplace. While this combination is in general undecidable, many restrictions of Presburger arithmetic with uninterpreted functions have been shown to be decidable. Stump et al. [189] present an algorithm that first eliminates the array update expressions from the formula by identifying matching writes. The resulting formula can be decided with an EUF decision procedure (Chap. 3).

The array property fragment that we used in this chapter was identified by Bradley, Manna and Sipma [31]. The idea of computing “sufficiently large” sets of instantiation values is also used in other procedures. For instance, Ghilardi et al. computed such sets separately for the indices and array elements [84].

## 7.6 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$a[i]$	The element with index $i$ of an array $a$	171
$a\{i \leftarrow e\}$	The array $a$ , where the element with index $i$ has been replaced by $e$	171
$\mathcal{I}(\phi)$	Index set	176

## Pointer Logic

### 8.1 Introduction

#### 8.1.1 Pointers and Their Applications

This chapter introduces a theory for reasoning about programs that use pointers, and describes decision procedures for it. We assume that the reader is familiar with pointers and their use in programming languages.

A **pointer** is a program variable whose sole purpose is to refer to some other program construct. This other construct could be a variable, a procedure or label, or yet another pointer. Among other things, pointers allow a piece of code to operate on different sets of data, which avoids inefficient copying of data.

As an example, consider a program that maintains two arrays of integers, named A and B, and that both arrays need to be sorted at some point within the program. Without pointers, the programmer needs to maintain two implementations of the sorting algorithm, one for A and one for B. Using pointers, a single implementation of sorting is implemented as a procedure that accepts a pointer to the first element of an array as an argument. It is called twice, with the addresses of A and B, respectively, as the argument.

As pointers are a common source of programming errors, most modern programming languages try to offer alternatives, e.g., in the form of references or abstract data containers. Nevertheless, low-level programming languages with explicit pointers are still frequently used, for example for embedded systems or operating systems.

The implementation of pointers relies on the fact that the memory cells of a computer have *addresses*, i.e., each cell has a unique number. The value of a pointer is then nothing but such a number. The way the memory cells are addressed is captured by the concept of the **memory model** of the architecture that executes the program.

**Definition 8.1 (memory model).** *A memory model describes the assumptions that are made about the way memory cells are addressed. We assume*



$M, A$ 

that the architecture provides a continuous, uniform address space, i.e., the set of addresses  $A$  is a subinterval of the integers  $\{0, \dots, N - 1\}$ . Each address corresponds to a memory cell that is able to store one data word. The set of data words is denoted by  $D$ . A **memory valuation**  $M : A \rightarrow D$  is a mapping from a set of addresses  $A$  into the domain  $D$  of data words.

 $D$ 

A variable may require more than one data word to be stored in memory. For example, this is the case when the variable is of type struct, array, or double-precision floating-point. Let  $\sigma(v)$  with  $v \in V$  denote the size (in data words) of  $v$ .

 $\sigma$ 

The compiler assigns a particular memory location (address) to each global, and thus, static variable.<sup>1</sup> This mapping is called the **memory layout**, and is formalized as follows. Let  $V$  denote the set of variables.

 $V$  $L$ 

**Definition 8.2 (memory layout).** A memory layout  $L : V \rightarrow A$  is a mapping from each variable  $v \in V$  to an address  $a \in A$ . The address of  $v$  is also called the memory location of  $v$ .

The memory locations of the statically allocated variables are usually assigned such that they are *nonoverlapping* (we explain later on how to model dynamically allocated data structures). Note that the memory layout is not necessarily continuous, i.e., the compilers may generate a layout that contains “holes”.<sup>2</sup>

**Example 8.3.** Figure 8.1 illustrates a memory layout for a fragment of an ANSI-C program. The program has six objects, which are named `var_a`, `var_b`, `var_c`, `s`, `array`, and `p`. The first five objects either are integer variables or are composed of integer variables. The object named `p` is a pointer variable, which we assume to be as wide as an integer.<sup>3</sup> The program initializes `p` to the address of the variable `var_c`, which is denoted by `&var_c`. Besides the variable definitions, the program also has a function `main()`, which sets the value of the variable pointed to by `p` to 100. ▀

### 8.1.2 Dynamic Memory Allocation

Pointers also enable the creation of *dynamic data structures*. Dynamic data structures rely on an area of memory that is designated for use by objects that

<sup>1</sup> Statically allocated variables are variables that are allocated space during the entire run time of the program. In contrast, the addresses of dynamically allocated data such as local variables or data on the heap are determined at run time once the object has been created.

<sup>2</sup> A possible reason for such holes is the need for proper alignment. As an example, many 64-bit architectures are unable to read double-precision floating-point values from addresses that are not a multiple of 8.


<sup>3</sup> This is not always the case; for example, in the x86 16-bit architecture, integers have 16 bits, whereas pointers are 32 bits wide. In some 64-bit architectures, integers have 32 bits, whereas pointers have 64 bits.

```

int var_a, var_b, var_c;
struct { int x; int y; } S;
int array[4];
int *p = &var_c;

int main() {
    *p=100;
}

```



**Fig. 8.1.** A fragment of an ANSI-C program and a possible memory layout for it

### Aside: Pointers and References in Object-Oriented Programming

Separation of data and algorithms is promoted by the concept of *object-oriented programming* (OOP). In modern programming languages such as Java and C++, the explicit use of pointer variables is deprecated. Instead, the procedures that are associated with an object (the *methods*) implicitly receive a pointer to the data members (the *fields*) of the object instance as an argument. In C++, the pointer is accessible using the keyword `this`. All accesses to the data members are performed indirectly by means of the `this` pointer variable.

**References**, just like pointers, are program variables that refer to a variable or object. The difference between references and pointers is often only syntactic. As an example, the fact that dereferencing is performed is usually hidden. In program analysis, references can be treated just as pointers.

are created at the run time of the program. A run time library maintains a list of the memory regions that are unused. A function, which is part of this library, allocates a region of given size and returns a pointer to the beginning (lowest address) of the region. The memory layout therefore *changes during the run time of the program*. Memory allocation may be performed an unbounded number of times (provided enough space is deallocated as well), and thus, there is no bound on the number of objects that a program can generate.

The function that performs the allocation is called `malloc()` in C, and is provided as an operator called `new` in C++, C#, and Java. In either case, the size of the region that is requested is passed as an argument. In order to reuse memory occupied by data structures that are no longer needed, C programmers call `free`, C++ programmers use `delete`, while Java and C# provide an automatic garbage collection mechanism. The **lifetime** of a dynamic object is the time between its allocation and its deallocation.

### 8.1.3 Analysis of Programs with Pointers

All but trivial programs rely on pointers or references in order to separate between data and algorithms. Decision procedures that are used for program analysis therefore often need to include reasoning about pointers.

As a simple example, consider the following program fragment, which computes the sum of an array of size 10:

```
void f(int *sum) {
    *sum = 0;

    for(i=0; i<10; i++)
        *sum = *sum + array[i];
}
```

The sum is stored in an integer variable that is pointed to by a pointer called `sum`. Any analysis method that aims at validating the correctness of this fragment has to take the value of the pointer into account. In particular, the program is likely to fail if the address held by `sum` is equal to the address of `i`. In this case, we say that `*sum` is an **alias** for `i`. Aliasing that is not anticipated by the programmer is a common source of problems.

The use of pointers gives rise to program properties that are of high interest. It is well known that many programs fail owing to incorrect use of pointer variables. A very common problem in programs is dereferencing of pointer variables that do not point to a proper object. The value 0 is typically reserved as a designated NULL pointer. It is guaranteed that no object, either statically or dynamically allocated, has this address. This value can therefore be used to indicate special cases, for example the end of a linked list. However, if such a pointer is – by mistake – dereferenced, modern architectures typically generate an exception, which terminates the program.

Programming languages that offer explicit deallocation face another problem. In the following program fragment, an array-type object is allocated and deallocated:

```
int *p, *q;

p = new int[10];
q = &p[3];
delete p;
*q = 2;
```

Note that the address of the fourth element of the array is stored in `q`, and that this pointer is dereferenced after the deallocation of the array. In a variant of the program above, the library that manages the dynamically allocated memory may have reassigned the space used for the array by that time, and thus another object might be overwritten by writing to `*q`. Such errors are hard to reproduce, as they depend on the exact memory layout of the architecture.

They often remain undetected despite extensive testing. The detection of such errors is therefore an important application for static program analysis tools.

### Aside: Alias Analysis

Alias analysis has a significant role in pointer-related reasoning about software, such as the analysis performed by optimizing compilers. Alias analysis may be performed at various levels of precision. For example, alias analysis may be field sensitive or insensitive, interprocedural or intra-procedural, and may or may not be sensitive to the control flow. Alias analysis is a special case of *static analysis*, and is typically performed as a *may-analysis*, that is, it determines the set of variables that a given pointer *may* point to – this is called the “points-to” set. In other words, variables that are *not* in this set cannot be pointed to by this pointer. For example, given an instruction such as

```
*p=0;
```

may-analysis permits us to conclude that any variable that is *not* in the points-to set of `p` is also not modified by this assignment. In the case of an optimizing compiler, this permits us to determine the set of variables that can be cached safely in processor registers.

Alias analysis is performed by maintaining a points-to set for each pointer (and, if desired, for each program location), and updating these sets according to the program statements. The algorithm terminates once the sets have saturated, i.e., do not change anymore.

As an example, consider a control-flow-insensitive analysis of a program with three statements:

```
p=q;  
q=&i;  
p=&j;
```

The points-to sets of `p` and `q` are initially empty. Processing the first statement results in no change. The second statement adds `i` to the points-to set of `q`, and the third adds `j` to the points-to set of `p`. Owing to the first statement, the set of `q` is added to that of `p` and, thereafter, the two sets are saturated.

## 8.2 A Simple Pointer Logic

### 8.2.1 Syntax

There are many variants of pointer logic, each with a different syntax and meaning. The more complex ones are often undecidable. We define a simple logic here, with the goal of making the problem of deciding formulas in this logic easier to solve.

**Definition 8.4 (pointer logic).** *The syntax of a formula in pointer logic is defined by the following rules:*

$$\begin{aligned}
 \text{formula} &: \text{formula} \wedge \text{formula} \mid \neg \text{formula} \mid (\text{formula}) \mid \text{atom} \\
 \text{atom} &: \text{pointer} = \text{pointer} \mid \text{term} = \text{term} \mid \\
 &\quad \text{pointer} < \text{pointer} \mid \text{term} < \text{term} \\
 \text{pointer} &: \text{pointer-identifier} \mid \text{pointer} + \text{term} \mid (\text{pointer}) \mid \\
 &\quad \& \text{identifier} \mid \& * \text{pointer} \mid * \text{pointer} \mid \text{NULL} \\
 \text{term} &: \text{identifier} \mid * \text{pointer} \mid \text{term op term} \mid (\text{term}) \mid \\
 &\quad \text{integer-constant} \mid \text{identifier} [ \text{term} ] \\
 \text{op} &: + \mid -
 \end{aligned}$$

The variables represented by *pointer-identifier* are assumed to be of pointer type, whereas the variables represented by *identifier* are assumed to be integers or an array of integers.<sup>4</sup> Note that the grammar allows pointer arithmetic, whereas it prohibits a direct conversion of an integer into a pointer or vice versa. This is motivated by the fact that the conversion of a pointer to an integer may actually fail in a number of architectures, owing to the fact that pointers are wider than the standard integer type.<sup>5</sup>

**Example 8.5.** Let  $p, q$  denote pointer identifiers, and let  $i, j$  denote integer identifiers. The following expressions are well-formed according to the grammar above:

- $*(p + i) = 1$ ,
- $*(p + *p) = 0$ ,
- $p = q \wedge *p = 5$ ,
- $****p = 1$ ,
- $p < q$ .

The following expressions are not permitted by the grammar:

- $p + i$ ,
- $p = i$ ,
- $*(p + q)$ ,
- $*1 = 1$ ,
- $p < i$ . ▀

Note that the grammar above encompasses all of integer linear arithmetic (Chap. 5) and also a fragment of array logic (Chap. 7). In practice, a logic for pointers is typically combined with a logic for the program expressions, such as bit-vector arithmetic.

<sup>4</sup> The syntax is clearly inspired by that of ANSI-C; Note, however, that we deviate from the ANSI-C syntax in a few points. As an example, in ANSI-C, an array identifier is synonymous with its address.

<sup>5</sup> Much as in C/C++, an indirect conversion by means of the dereferencing operator is still possible.

### 8.2.2 Semantics

There are numerous ways to assign a meaning to the expressions defined above. We define the semantics by referring to a specific memory layout  $L$  (Definition 8.2) and a specific memory valuation  $M$  (Definition 8.1), that is, pointer logic formulas are predicates on  $M, L$  pairs. The definition uses a reduction to integer arithmetic and array logic, and thus we treat  $M$  and  $L$  as array types. We also assume that  $D$  (the set of data words) is contained in the set of integers.

**Definition 8.6 (semantics of pointer logic).** *As before let  $L$  denote a memory layout and let  $M$  denote a valuation of the memory. Let  $\mathcal{L}_P$  denote the set of pointer logic expressions, and let  $\mathcal{L}_D$  denote the set of expressions permitted by the logic for the data words. We define a meaning for  $e \in \mathcal{L}_P$  using the function  $\llbracket \cdot \rrbracket : \mathcal{L}_P \rightarrow \mathcal{L}_D$ . The function  $\llbracket e \rrbracket$  is defined recursively as given in Fig. 8.2. The expression  $e \in \mathcal{L}_P$  is valid if and only if  $\llbracket e \rrbracket$  is valid.*

$\llbracket f_1 \wedge f_2 \rrbracket$	$\doteq$	$\llbracket f_1 \rrbracket \wedge \llbracket f_2 \rrbracket$	
$\llbracket \neg f \rrbracket$	$\doteq$	$\neg \llbracket f \rrbracket$	
$\llbracket p_1 = p_2 \rrbracket$	$\doteq$	$\llbracket p_1 \rrbracket = \llbracket p_2 \rrbracket$	where $p_1, p_2$ are pointer expressions
$\llbracket p_1 < p_2 \rrbracket$	$\doteq$	$\llbracket p_1 \rrbracket < \llbracket p_2 \rrbracket$	where $p_1, p_2$ are pointer expressions
$\llbracket t_1 = t_2 \rrbracket$	$\doteq$	$\llbracket t_1 \rrbracket = \llbracket t_2 \rrbracket$	where $t_1, t_2$ are terms
$\llbracket t_1 < t_2 \rrbracket$	$\doteq$	$\llbracket t_1 \rrbracket < \llbracket t_2 \rrbracket$	where $t_1, t_2$ are terms
$\llbracket p \rrbracket$	$\doteq$	$M[L[p]]$	where $p$ is a pointer identifier
$\llbracket p + t \rrbracket$	$\doteq$	$\llbracket p \rrbracket + \llbracket t \rrbracket$	where $p$ is a pointer expression, and $t$ is a term
$\llbracket \&v \rrbracket$	$\doteq$	$L[v]$	where $v \in V$ variable
$\llbracket \& * p \rrbracket$	$\doteq$	$\llbracket p \rrbracket$	where $p$ is a pointer expression
$\llbracket \text{NULL} \rrbracket$	$\doteq$	$0$	
$\llbracket v \rrbracket$	$\doteq$	$M[L[v]]$	where $v \in V$ is a variable
$\llbracket *p \rrbracket$	$\doteq$	$M[\llbracket p \rrbracket]$	where $p$ is a pointer expression
$\llbracket t_1 \text{ op } t_2 \rrbracket$	$\doteq$	$\llbracket t_1 \rrbracket \text{ op } \llbracket t_2 \rrbracket$	where $t_1, t_2$ are terms
$\llbracket c \rrbracket$	$\doteq$	$c$	where $c$ is an integer constant
$\llbracket v[t] \rrbracket$	$\doteq$	$M[L[v] + \llbracket t \rrbracket]$	where $v$ is an array identifier, and $t$ is a term

**Fig. 8.2.** Semantics of pointer expressions

Observe that a pointer  $p$  points to a variable  $x$  if  $M[L[p]] = L[x]$ , that is, the value of  $p$  is equal to the address of  $x$ . As a shorthand, we write  $p \hookrightarrow z$  to mean that  $p$  points to some memory cell such that  $*p = z$ . Observe also that the meaning of pointer arithmetic, for example  $p + i$ , does not depend on the type of the object that  $p$  points to.<sup>6</sup>

$p \hookrightarrow z$

<sup>6</sup> In contrast, the semantics of ANSI-C requires that an integer that is added to a pointer  $p$  is multiplied by the size of the type that  $p$  points to.

**Example 8.7.** Consider the following expression, where  $a$  is an array identifier:

$$*((\&a) + 1) = a[1] . \quad (8.1)$$

The semantic definition of (8.1) expands as follows:

$$\llbracket *((\&a) + 1) = a[1] \rrbracket \iff \llbracket *((\&a) + 1) \rrbracket = \llbracket a[1] \rrbracket \quad (8.2)$$

$$\iff M[\llbracket (\&a) + 1 \rrbracket] = M[L[a] + \llbracket 1 \rrbracket] \quad (8.3)$$

$$\iff M[\llbracket \&a \rrbracket + \llbracket 1 \rrbracket] = M[L[a] + 1] \quad (8.4)$$

$$\iff M[L[a] + 1] = M[L[a] + 1] \quad (8.5)$$

Equation (8.5) is obviously valid, and thus, so is (8.1). Note that the translated formula must evaluate to true for any  $L$  and  $M$  and, thus, the following formula is not valid:

$$*p = 1 \implies x = 1 . \quad (8.6)$$

For  $p \neq \&x$ , this formula evaluates to FALSE. ■

### 8.2.3 Axiomatization of the Memory Model

Formulas in pointer logic may exploit assumptions made about the memory model. The set of these assumptions depends highly on the architecture. Here, we formalize properties that most architectures comply with, and thus that many programs rely on.

On most architectures, the following two formulas are valid, and hence can be safely assumed by programmers:

$$\&x \neq \text{NULL} , \quad (8.7)$$

$$\&x \neq \&y . \quad (8.8)$$

Equation (8.7) translates into  $L[x] \neq 0$  and relies on the fact that no object has address 0. Equation (8.8) relies on the fact that the memory layout assigns nonoverlapping addresses to the objects. We define a series of *memory model axioms* in order to formalize these properties.

**Memory Model Axiom 1 (“No object has address 0”)** *The fact “no object has address 0” is easily formalized:*<sup>7</sup>

$$\forall v \in V. L[v] \neq 0 . \quad (8.9)$$

---

<sup>7</sup> Note that the ANSI-C standard does not actually guarantee that the symbolic constant NULL is represented by a bit vector consisting of zeros; however, it guarantees that the NULL pointer compares to the integer zero and can be obtained by converting the integer zero to a pointer type.

The easiest way to ensure that (8.8) is valid is to assume that  $\forall v_1, v_2 \in V. v_1 \neq v_2 \implies L[v_1] \neq L[v_2]$ . However, this assumption is often not strong enough, as objects with size greater or equal to two may still overlap. We therefore assume the following two conditions, which together are stronger.

**Memory Model Axiom 2 (“Objects have size at least one”)** *The fact “an object has size at least one” is easily captured by:*

$$\forall v \in V. \sigma(v) \geq 1. \quad (8.10)$$

**Memory Model Axiom 3 (“Objects do not overlap”)** *Different objects do not share any addresses:*

$$\forall v_1, v_2 \in V. v_1 \neq v_2 \implies \{L[v_1], \dots, L[v_1] + \sigma(v_1) - 1\} \cap \{L[v_2], \dots, L[v_2] + \sigma(v_2) - 1\} = \emptyset. \quad (8.11)$$

Program analysis tools that are applied to code that relies on additional, architecture-specific guarantees, may require a larger set of memory model axioms. Examples are *byte ordering* and *endianness*, and specific assumptions about *alignment*.

## 8.2.4 Adding Structure Types

Structure types are a convenient way to implement data structures. Structure types can be added to our pointer logic as a purely syntactic extension, as we shall soon see. We assume that the fields of the structure types are named, and write  $s.f$  to denote the value of the field  $f$  in the structure  $s$ .

Formally, we can view structure types as “syntactic sugar” for array types, and record the following shorthands. Each field of the structure is assigned a unique **offset**. Let  $o(f)$  denote the offset of field  $f$ . We then define the meaning of  $s.f$  as follows:

$$s.f \doteq *((\&s) + o(f)). \quad (8.12)$$

For convenience, we introduce two additional shorthands. Following the PASCAL and ANSI-C syntax, we write  $p \rightarrow f$  for  $(*p).f$  (this shorthand is not to be confused with logical implication or with  $p \hookrightarrow a$ ). Adopting some notation from separation logic (see the aside on separation logic), we also extend the  $p \hookrightarrow a$  notation by introducing  $p \hookrightarrow a, b, c, \dots$  as a shorthand for

$$\begin{aligned} *(p + 0) &= a \wedge \\ *(p + 1) &= b \wedge \\ *(p + 2) &= c \dots \end{aligned} \quad (8.13)$$

$$o(f)$$

$$s.f$$

$$p \rightarrow f$$




## 8.3 Modeling Heap-Allocated Data Structures

### 8.3.1 Lists

Heap-allocated data structures play an important role in programs, and are prone to pointer-related errors. We now illustrate how to model a number of commonly used data structures using pointer logic.

After the array, the simplest dynamically allocated data structure is the *linked list*. It is typically realized by means of a structure type that contains fields for a *next pointer* and the data that is to be stored in the list.

As an example, consider the following list. The first field is named  $a$  and is an ASCII character, serving as the “payload”, and the second field is named  $n$ , and is the pointer to the next element of the list. Following ANSI-C syntax, we use ‘x’ to denote the integer that represents the ASCII character “x”:



The list is terminated by a NULL pointer, which is denoted by “0” in the diagram above. A way of modeling this list is to use the following formula:

$$\begin{aligned}
 & p \hookrightarrow \text{'t'}, p_1 \\
 \wedge & p_1 \hookrightarrow \text{'e'}, p_2 \\
 \wedge & p_2 \hookrightarrow \text{'x'}, p_3 \\
 \wedge & p_3 \hookrightarrow \text{'t'}, \text{NULL} .
 \end{aligned} \tag{8.14}$$

This way of specifying lists is cumbersome, however. Therefore, disregarding the payload field, we first introduce a recursive shorthand for the  $i$ -th member of a list:<sup>8</sup>


$$\begin{aligned}
 \text{list-elem}(p, 0) & \doteq p , \\
 \text{list-elem}(p, i) & \doteq \text{list-elem}(p, i - 1) \rightarrow n \quad \text{for } i \geq 1 .
 \end{aligned} \tag{8.15}$$

list

We now define the shorthand  $\text{list}(p, l)$  to denote a predicate that is true if  $p$  points to a NULL-terminated acyclic list of length  $l$ :

$$\text{list}(p, l) \doteq \text{list-elem}(p, l) = \text{NULL} . \tag{8.16}$$

A linked list is *cyclic* if the pointer of the last element points to the first one:




Consider the following variant  $\text{my-list}(p, l)$ , intended to capture the fact that  $p$  points to such a cyclic list of length  $l \geq 1$ :

<sup>8</sup> Note that recursive definitions of this form are, in general, only embeddable into our pointer logic if the second argument is a constant.

$$\mathbf{my-list}(p, l) \doteq \mathbf{list-elem}(p, l) = p. \quad (8.17)$$

Does this definition capture the concept properly? The list in the diagram above satisfies  $\mathbf{my-list}(p, 4)$ . Unfortunately, the following list satisfies  $\mathbf{my-list}(p, 4)$  just as well:



This is due to the fact that our definition does not preclude *sharing* of elements of the list, despite the fact that we had certainly intended to specify that there are  $l$  disjoint list elements. Properties of this kind are often referred to as *separation properties*. A way to assert that the list elements are disjoint is to define a shorthand **overlap** as follows:

$$\mathbf{overlap}(p, q) \doteq p = q \vee p + 1 = q \vee p = q + 1. \quad (8.18)$$

This shorthand is then used to state that all list elements are pairwise disjoint:

$$\begin{aligned} \mathbf{list-disjoint}(p, 0) &\doteq \text{TRUE}, \\ \mathbf{list-disjoint}(p, l) &\doteq \mathbf{list-disjoint}(p, l - 1) \wedge \\ &\quad \forall 0 \leq i < l - 1. \neg \mathbf{overlap}(\mathbf{list-elem}(p, i), \mathbf{list-elem}(p, l - 1)). \end{aligned} \quad (8.19)$$

Note that the size of this formula grows quadratically in  $l$ . As separation properties are frequently needed, more concise notations have been developed for this concept, for example *separation logic* (see the aside on that topic). Separation logic can express such properties with formulas of linear size.

### 8.3.2 Trees

We can implement a *binary tree* by adding another pointer field to each element of the data structure (see Fig. 8.3). We denote the pointer to the left-hand child node by  $l$ , and the pointer to the right-hand child by  $r$ .

In order to illustrate a pointer logic formula for trees, consider the tree in Fig. 8.3, which has one integer  $x$  as payload. Observe that the integers are arranged in a particular fashion: the integer of the left-hand child of any node  $n$  is always smaller than the integer of the node  $n$  itself, whereas the integer of the right-hand child of node  $n$  is always larger than the integer of the node  $n$ . This property permits lookup of elements with a given integer value in time  $O(h)$ , where  $h$  is the height of the tree. The property can be formalized as follows:

$$\begin{aligned} (n.l \neq \text{NULL} \implies n.l \rightarrow x < n.x) \\ \wedge (n.r \neq \text{NULL} \implies n.r \rightarrow x > n.x). \end{aligned} \quad (8.22)$$

Unfortunately, (8.22) is not strong enough to imply lookup in time  $O(h)$ . For this, we need to establish the ordering over the integers of an *entire subtree*.

### Aside: Separation Logic

Theories for dynamic data structures are frequently used for proving that memory cells *do not alias*. While it is possible to model the statement that a given object does not alias with other objects with pairwise comparison, reasoning about such formulation scales poorly. It requires enumeration of all heap-allocated objects, which makes it difficult to reason about a program in a local manner.


John Reynolds' *separation logic* [165] addresses both problems by introducing a new binary operator “\*”, as in “ $P * Q$ ”, which is called a *separating conjunction*. The meaning of  $*$  is similar to the standard Boolean conjunction, i.e.,  $P \wedge Q$ , but it also asserts that  $P$  and  $Q$  reason about separate, nonoverlapping portions of the heap. As an example, consider the following variant of the `list` predicate:

$$\begin{aligned} \text{list}(p, 0) &\doteq p = \text{NULL} \\ \text{list}(p, l) &\doteq \exists q. p \hookrightarrow z, q \wedge \text{list}(q, l - 1) \quad \text{for } l \geq 1. \end{aligned} \quad (8.20)$$

Like our previous definition, the definition above suffers from the fact that some memory cells of the elements of the list might overlap. This can be mended by replacing the standard conjunction in the definition above by a separating conjunction:

$$\text{list}(p, l) \doteq \exists q. p \hookrightarrow z, q * \text{list}(q, l - 1). \quad (8.21)$$

This new list predicate also asserts that the memory cells of all list elements are pairwise disjoint. Separation logic, in its generic form, is not decidable, but a variety of decidable fragments have been identified.



**Fig. 8.3.** A binary tree that represents a set of integers

We define a predicate **tree-reach**( $p, q$ ), which holds if  $q$  is reachable from  $p$  in one step:

$$\begin{aligned} \text{tree-reach}(p, q) \doteq & p \neq \text{NULL} \wedge q \neq \text{NULL} \wedge \\ & (p = q \vee p \rightarrow l = q \vee p \rightarrow r = q) . \end{aligned} \quad (8.23)$$

In order to obtain a predicate that holds if and only if  $q$  is reachable from  $p$  in any number of steps, we define the *transitive closure* of a given binary relation  $R$ .

**Definition 8.8 (transitive closure).** *Given a binary relation  $R$ , the transitive closure  $\text{TC}_R$  relates  $x$  and  $y$  if there are  $z_1, z_2, \dots, z_n$  such that*

$$xRz_1 \wedge z_1Rz_2 \wedge \dots \wedge z_nRy .$$

*Formally, transitive closure can be defined inductively as follows:*

$$\begin{aligned} \text{TC}_R^1(p, q) & \doteq R(p, q) , \\ \text{TC}_R^i(p, q) & \doteq \exists p'. \text{TC}_R^{i-1}(p, p') \wedge R(p', q) \\ \text{TC}(p, q) & \doteq \exists i. \text{TC}_R^i(p, q) . \end{aligned} \quad (8.24)$$

Using the transitive closure of our **tree-reach** relation, we obtain a new relation **tree-reach\***( $p, q$ ) that holds if and only if  $q$  is reachable from  $p$  in any number of steps:

$$\text{tree-reach}^*(p, q) \iff \text{TC}_{\text{tree-reach}}(p, q) . \quad (8.25)$$

Using **tree-reach\***, it is easy to strengthen (8.22) appropriately:

$$\begin{aligned} (\forall p. \text{tree-reach}^*(n.l, p) \implies p \rightarrow x < n.x) \\ \wedge (\forall p. \text{tree-reach}^*(n.r, p) \implies p \rightarrow x > n.x) . \end{aligned} \quad (8.26)$$

Unfortunately, the addition of the transitive closure operator can make even simple logics undecidable, and thus, while convenient for modeling, it is a burden for automated reasoning. We restrict the presentation below to decidable cases by considering only special cases.

## 8.4 A Decision Procedure

### 8.4.1 Applying the Semantic Translation

The semantic translation introduced in Sect. 8.2.2 not only assigns meaning to the pointer formulas, but also gives rise to a simple decision procedure. The formulas generated by this semantic translation contain array read operators and linear arithmetic over the type that is used for the indices. This may be the set of integers (Chap. 5) or the set of bit vectors (Chap. 6). It also

contains at least equalities over the type that is used to model the contents of the memory cells. We assume that this is the same type as the index type. As we have seen in Chap. 7, such a logic is decidable. Care has to be taken when extending the pointer logic with quantification, as array logic with arbitrary quantification is undecidable.

A straightforward decision procedure for pointer logic therefore first applies the semantic translation to a pointer formula  $\varphi$  to obtain a formula  $\varphi'$  in the combined logic of linear arithmetic over integers and arrays of integers. The formula  $\varphi'$  is then passed to the decision procedure for the combined logic. As the formulas  $\varphi$  and  $\varphi'$  are equisatisfiable (by definition), the result returned for  $\varphi'$  is also the correct result for  $\varphi$ .

**Example 8.9.** Consider the following pointer logic formula, where  $x$  is a variable, and  $p$  identifies a pointer:

$$p = \&x \wedge x = 1 \implies *p = 1. \quad (8.27)$$

The semantic definition of this formula expands as follows:

$$\begin{aligned} & \llbracket p = \&x \wedge x = 1 \implies *p = 1 \rrbracket \\ & \iff \llbracket p = \&x \rrbracket \wedge \llbracket x = 1 \rrbracket \implies \llbracket *p = 1 \rrbracket \\ & \iff \llbracket p \rrbracket = \llbracket \&x \rrbracket \wedge \llbracket x \rrbracket = 1 \implies \llbracket *p \rrbracket = 1 \\ & \iff M[L[p]] = L[x] \wedge M[L[x]] = 1 \implies M[M[L[p]]] = 1. \end{aligned} \quad (8.28)$$

A decision procedure for array logic and equality logic easily concludes that the formula above is valid, and thus, so is (8.27).

As an example of an invalid formula, consider

$$p \hookrightarrow x \implies p = \&x. \quad (8.29)$$


The semantic definition of this formula expands as follows:

$$\begin{aligned} & \llbracket p \hookrightarrow x \implies p = \&x \rrbracket \\ & \iff \llbracket p \hookrightarrow x \rrbracket \implies \llbracket p = \&x \rrbracket \\ & \iff \llbracket *p = x \rrbracket \implies \llbracket p \rrbracket = \llbracket \&x \rrbracket \\ & \iff \llbracket *p \rrbracket = \llbracket x \rrbracket \implies M[L[p]] = L[x] \\ & \iff M[M[L[p]]] = M[L[x]] \implies M[L[p]] = L[x] \end{aligned} \quad (8.30)$$

A counterexample to this formula is the following:

$$L[p] = 1, L[x] = 2, M[1] = 3, M[2] = 10, M[3] = 10. \quad (8.31)$$

The values of  $M$  and  $L$  in the counterexample are best illustrated with a picture:



□

### Applying the Memory Model Axioms

A formula may rely on one of the memory model axioms defined in Sect. 8.2.3. As an example, consider the following formula:

$$\sigma(x) = 2 \implies \&x \neq \&x + 1 . \quad (8.32)$$

The semantic translation yields:

$$\sigma(x) = 2 \implies L[y] \neq L[x] + 1 . \quad (8.33)$$

This formula can be shown to be valid by instantiating Memory Model Axiom 3. After instantiating  $v_1$  with  $x$  and  $v_2$  with  $y$ , we obtain

$$\{L[x], \dots, L[x] + \sigma(x) - 1\} \cap \{L[y], \dots, L[y] + \sigma(y) - 1\} = \emptyset . \quad (8.34)$$

We can transform the set expressions in (8.34) into linear arithmetic over the integers as follows:

$$(L[x] + \sigma(x) - 1 < L[y]) \vee (L[x] > L[y] + \sigma(y) - 1) . \quad (8.35)$$

Using  $\sigma(x) = 2$  and  $\sigma(y) \geq 1$  (Memory Model Axiom 2), we can conclude, furthermore, that

$$(L[x] + 1 < L[y]) \vee (L[x] > L[y]) . \quad (8.36)$$

Equation (8.36) is strong enough to imply  $L[y] \neq L[x] + 1$ , which proves that Eq. (8.32) is valid.

#### 8.4.2 Pure Variables

The semantic translation of a pointer formula results in a formula that we can decide using the procedures described in this book. However, the semantic translation down to memory valuations places an undue burden on the underlying decision procedure, as illustrated by the following example (symmetry of equality):

$$\llbracket x = y \implies y = x \rrbracket \quad (8.37)$$

$$\iff \llbracket x = y \rrbracket \implies \llbracket y = x \rrbracket \quad (8.38)$$

$$\iff M[L[x]] = M[L[y]] \implies M[L[y]] = M[L[x]] . \quad (8.39)$$

A decision procedure for array logic and equality logic is certainly able to deduce that (8.39) is valid. Nevertheless, the steps required for solving (8.39) obviously exceed the effort required to decide

$$x = y \implies y = x . \quad (8.40)$$

In particular, the semantic translation does not exploit the fact that  $x$  and  $y$  do not actually interact with any pointers. A straightforward optimization is therefore the following: if the address of a variable  $x$  is not referred to, we translate it to a new variable  $\Upsilon_x$  instead of  $M[L[x]]$ . A formalization of this idea requires the following definition:

**Definition 8.10 (pure variables).** Given a formula  $\varphi$  with a set of variables  $V$ , let  $\mathcal{P}(\varphi) \subseteq V$  denote the subset of  $\varphi$ 's variables that are not used within an argument of the “&” operator within  $\varphi$ . These variables are called pure.

As an example,  $\mathcal{P}(\&x = y)$  is  $\{y\}$ . We now define a new translation function  $\llbracket \cdot \rrbracket^{\mathcal{P}}$ . The definition of  $\llbracket e \rrbracket^{\mathcal{P}}$  is identical to the definition of  $\llbracket e \rrbracket$  unless  $e$  denotes a variable in  $\mathcal{P}(\varphi)$ . The new definition is:

$$\begin{aligned} \llbracket v \rrbracket^{\mathcal{P}} &\doteq \Upsilon_v && \text{for } v \in \mathcal{P}(\varphi) \\ \llbracket v \rrbracket^{\mathcal{P}} &\doteq M[L[v]] && \text{for } v \in V \setminus \mathcal{P}(\varphi) \end{aligned}$$

**Theorem 8.11.** The translation using pure variables is equisatisfiable with the semantic translation:

$$\llbracket \varphi \rrbracket^{\mathcal{P}} \iff \llbracket \varphi \rrbracket .$$

**Example 8.12.** Equation (8.38) is now translated as follows without referring to a memory valuation, and thus no longer burdens the decision procedure for array logic:

$$\llbracket x = y \implies y = x \rrbracket^{\mathcal{P}} \tag{8.41}$$

$$\iff \llbracket x = y \implies y = x \rrbracket^{\mathcal{P}} \tag{8.42}$$

$$\iff \llbracket x = y \rrbracket^{\mathcal{P}} \implies \llbracket y = x \rrbracket^{\mathcal{P}} \tag{8.43}$$

$$\iff \Upsilon_x = \Upsilon_y \implies \Upsilon_y = \Upsilon_x . \tag{8.44}$$

▀

### 8.4.3 Partitioning the Memory

The translation procedure can be optimized further using the following observation: the run time of a decision procedure for array logic depends on the number of different expressions that are used to index a particular array (see Chap. 7). As an example, consider the pointer logic formula

$$*p = 1 \wedge *q = 1 , \tag{8.45}$$

which – using our optimized translation – is reduced to

$$M[\Upsilon_p] = 1 \wedge M[\Upsilon_q] = 1 . \tag{8.46}$$

The pointers  $p$  and  $q$  might alias, but there is no reason why they have to. Without loss of generality, we can therefore safely assume that they do not alias and, thus, we partition  $M$  into  $M_1$  and  $M_2$ :

$$M_1[\Upsilon_p] = 1 \wedge M_2[\Upsilon_q] = 1 . \tag{8.47}$$

While this has increased the number of array variables, the number of different indices *per array* has decreased. Typically, this improves the performance of a decision procedure for array logic.

This transformation cannot always be applied, which is illustrated by the following example:

$$p = q \implies *p = *q . \quad (8.48)$$

This formula is obviously valid, but if we partition as before, the translated formula is no longer valid:

$$\Upsilon_p = \Upsilon_q \implies M_1[\Upsilon_p] = M_2[\Upsilon_q] . \quad (8.49)$$

Unfortunately, deciding if the optimization is applicable is in general as hard as deciding  $\varphi$  itself. We therefore settle for an approximation based on a syntactic test. This approximation is conservative, i.e., sound, while it may not result in the best partitioning that is possible in theory.

**Definition 8.13.** *We say that two pointer expressions  $p$  and  $q$  are related directly by a formula  $\varphi$  if both  $p$  and  $q$  are used inside the same relational expression in  $\varphi$  and that the expressions are related transitively if there is a pointer expression  $p'$  that relates to  $p$  and relates to  $q$ . We write  $p \approx q$  if  $p$  and  $q$  are related directly or transitively.*

$$p \approx q$$

The relation  $\approx$  induces a partitioning of the pointer expressions in  $\varphi$ . We number these partitions  $1, \dots, n$ . Let  $I(p) \in \{1, \dots, n\}$  denote the index of the partition that  $p$  is in. We now define a new translation  $\llbracket \cdot \rrbracket^\approx$ , in which we use a separate memory valuation  $M_{I(p)}$  when  $p$  is dereferenced. The definition of  $\llbracket e \rrbracket^\approx$  is identical to the definition of  $\llbracket e \rrbracket^P$  unless  $e$  is a dereferencing expression. In this case, we use the following definition:

$$\llbracket *p \rrbracket^\approx \doteq M_{I(p)}(\llbracket p \rrbracket^\approx) .$$

**Theorem 8.14.** *Translation using memory partitioning results in a formula that is equisatisfiable with the result of the semantic translation:*

$$\exists \alpha_1. \alpha_1 \models \llbracket \varphi \rrbracket^\approx \iff \exists \alpha_2. \alpha_2 \models \llbracket \varphi \rrbracket .$$

Note that the theorem relies on the fact that our grammar does not permit explicit restrictions on the memory layout  $L$ . The theorem no longer holds as soon as this restriction is lifted (see Problem 8.5).

## 8.5 Rule-Based Decision Procedures

With pointer logics expressive enough to model interesting data structures, one often settles for incomplete, rule-based procedures. The basic idea of such procedures is to define a fragment of pointer logic enriched with predicates for specific types of data structures (e.g., lists or trees) together with a set of proof rules that are sufficient to prove a wide range of verification conditions that arise in practice. The soundness of these proof rules is usually shown with respect to the definitions of the predicates, which implies soundness of the decision procedure. There are only a few known proof systems that are provably complete.



### 8.5.1 A Reachability Predicate for Linked Structures

As a simple example of this approach, we present a variant of a calculus for reachability predicates introduced by Greg Nelson [135]. Further rule-based reasoning systems are discussed in the bibliographic notes at the end of this chapter.

We first generalize the **list-*elem*** shorthand used before for specifying linked lists by parameterizing it with the name of the field that holds the pointer to the “next” element. Suppose that  $f$  is a field of a structure and holds a pointer. The shorthand  $\mathbf{follow}_n^f(q)$  stands for the pointer that is obtained by starting from  $q$  and following the field  $f$ ,  $n$  times:

$$\begin{aligned} \mathbf{follow}_0^f(p) &\doteq p \\ \mathbf{follow}_n^f(p) &\doteq \mathbf{follow}_{n-1}^f(p) \rightarrow f . \end{aligned} \tag{8.50}$$

If  $\mathbf{follow}_n^f(p) = q$  holds, then  $q$  is reachable in  $n$  steps from  $p$  by following  $f$ . We say that  $q$  is reachable from  $p$  by following  $f$  if there exists such  $n$ . Using this shorthand, we enrich the logic with just a single predicate for list-like data structures, denoted by

$$p \xrightarrow[x]{f} q , \tag{8.51}$$

which is called a **reachability predicate**. It is read as “ $q$  is reachable from  $p$  following  $f$ , while avoiding  $x$ ”. It holds if two conditions are fulfilled:

1. There exists some  $n$  such that  $q$  is reachable from  $p$  by following  $f$   $n$  times.
2.  $x$  is not reachable in fewer than  $n$  steps from  $p$  following  $f$ .

This can be formalized using  $\mathbf{follow}()$  as follows:

$$p \xrightarrow[x]{f} q \iff \exists n. (\mathbf{follow}_n^f(p) = q \wedge \forall m < n. \mathbf{follow}_m^f(p) \neq x) . \tag{8.52}$$

We say that a formula is a **reachability predicate formula** if it contains the reachability predicate.

**Example 8.15.** Consider the following software verification problem. The following program fragment iterates over an acyclic list and searches for a list entry with payload  $a$ :

```
struct S { struct S *nxt; int payload; } *list;
...
bool find(int a) {
    for(struct S *p=list; p!=0; p=p->nxt)
        if(p->payload==a) return true;
    return false;
}
```

We can specify the correctness of the result returned by this procedure using the following formula:

$$\text{find}(a) \iff \exists p'. (\text{list} \xrightarrow[0]{\text{next}} p' \wedge p' \rightarrow \text{payload} = a) . \quad (8.53)$$

Thus,  $\text{find}(a)$  is true if the following conditions hold:

1. There is a list element that is reachable from  $\text{list}$  by following  $\text{next}$  without passing through a NULL pointer.
2. The payload of this list element is equal to  $a$ .

We annotate the beginning of the loop body in the program above with the following loop invariant, denoted by **INV**:

$$\text{INV} := \text{list} \xrightarrow[0]{\text{next}} p \wedge (\forall q \neq p. \text{list} \xrightarrow[p]{\text{next}} q \implies q \rightarrow \text{payload} \neq a) . \quad (8.54)$$

Informally, we make the following argument: first, we show that the program maintains the loop invariant **INV**; then, we show that **INV** implies our property.

Formally, this is shown by means of four **verification conditions**. The validity of all of these verification conditions implies the property. We use the notation  $e[x/y]$  to denote the expression  $e$  in which  $x$  is replaced by  $y$ .

$$\text{IND-BASE} := p = \text{list} \implies \text{INV} \quad (8.55)$$

$$\text{IND-STEP} := (\text{INV} \wedge p \rightarrow \text{payload} \neq a) \implies \text{INV}[p/p \rightarrow \text{next}] \quad (8.56)$$

$$\text{VC-P1} := (\text{INV} \wedge p \rightarrow \text{payload} = a) \quad (8.57)$$

$$\implies \exists p'. (\text{list} \xrightarrow[0]{\text{next}} p' \wedge p' \rightarrow \text{payload} = a)$$

$$\text{VC-P2} := (\text{INV} \wedge p = 0) \implies \neg \exists p'. (\text{list} \xrightarrow[0]{\text{next}} p' \wedge p' \rightarrow \text{payload} = a) \quad (8.58)$$

The first verification condition, **IND-BASE**, corresponds to the induction base of the inductive proof. It states that **INV** holds upon entering the loop, because at that point  $p = \text{list}$ . The formula **IND-STEP** corresponds to the induction step: it states that the loop-invariant is maintained if another loop iteration is executed (i.e.,  $p \rightarrow \text{payload} \neq a$ ).

The formulas **VC-P1** and **VC-P2** correspond to the two cases of leaving the  $\text{find}$  function: **VC-P1** establishes the property if **TRUE** is returned, and **VC-P2** establishes the property if **FALSE** is returned. Proving these verification conditions therefore shows that the program satisfies the required property. ▀

### 8.5.2 Deciding Reachability Predicate Formulas

As before, we can simply expand the definition above and obtain a semantic reduction. As an example, consider the verification condition labeled **IND-BASE** in Sect. 8.5.1:

$$p = list \implies \text{INV} \quad (8.59)$$

$$\iff p = list \implies list \xrightarrow[nxt]{0} p \wedge \forall q \neq p. list \xrightarrow[nxt]{p} q \implies q \rightarrow \text{payload} \neq a \quad (8.60)$$

$$\iff list \xrightarrow[nxt]{0} list \wedge \forall q \neq list. (list \xrightarrow[nxt]{list} q \implies q \rightarrow \text{payload} \neq a) \quad (8.61)$$

$$\begin{aligned} \iff (\exists n. \text{follow}_n^{nxt}(list) = list \wedge \forall m < n. \text{follow}_m^{nxt}(list) \neq list) \wedge \\ (\forall q \neq list. ((\exists n. \text{follow}_n^{nxt}(list) = q \wedge \forall m < n. \text{follow}_m^{nxt}(list) \neq list) \\ \implies q \rightarrow \text{payload} \neq a)) . \end{aligned} \quad (8.62)$$

Equation (8.62) is argued to be valid as follows. In the first conjunction, instantiate  $n$  with 0. In the second conjunct, observe that  $q \neq list$ , and thus any  $n$  satisfying  $\exists n. \text{follow}_n^{nxt}(list) = q$  must be greater than 0. Finally, observe that  $\text{follow}_m^{nxt}(list) \neq list$  is invalid for  $m = 0$ , and thus the left-hand side of the implication is FALSE.

However, note that the formulas above contain many existential and universal quantifiers over natural numbers and pointers. Applying the semantic reduction therefore does not result in a formula that is in the array property fragment defined in Chap. 7. Thus, the decidability result shown in this chapter does not apply here. How can such complex reachability predicate formulas be solved?

## Using Rules

In such situations, the following technique is frequently applied: *rules* are derived from the semantic definition of the predicate, and then they are applied to simplify the formula.

$$p \xrightarrow{f}{x} q \iff (p = q \vee (p \neq x \wedge p \rightarrow f \xrightarrow{f}{x} q)) \quad (A1)$$

$$(p \xrightarrow{f}{x} q \wedge q \xrightarrow{f}{x} r) \implies p \xrightarrow{f}{x} r \quad (A2)$$

$$p \xrightarrow{f}{x} q \implies p \xrightarrow{f}{q} q \quad (A3)$$

$$(p \xrightarrow{f}{y} x \wedge p \xrightarrow{f}{z} y) \implies p \xrightarrow{f}{z} x \quad (A4)$$

$$(p \xrightarrow{f}{x} x \vee p \xrightarrow{f}{y} y) \implies (p \xrightarrow{f}{y} x \vee p \xrightarrow{f}{x} y) \quad (A5)$$

$$(p \xrightarrow{f}{y} x \wedge p \xrightarrow{f}{z} y) \implies x \xrightarrow{f}{z} y \quad (A6)$$

$$p \rightarrow f \xrightarrow{f}{q} q \iff p \rightarrow f \xrightarrow{f}{p} q \quad (A7)$$

**Fig. 8.4.** Rules for the reachability predicate

The rules provided in [135] for our reachability predicate are given in Fig. 8.4. The first rule (A1) corresponds to a program fragment that follows field  $f$  once. If  $q$  is reachable from  $p$ , avoiding  $x$ , then either  $p = q$  (we are already there) or  $p \neq x$ , and we can follow  $f$  from  $p$  to get to a node from which  $q$  is reachable, avoiding  $x$ . We now prove the correctness of this rule.

*Proof.* We first expand the definition of our reachability predicate:

$$p \xrightarrow[f]{x} q \iff \exists n. (\mathbf{follow}_n^f(p) = q \wedge \forall m < n. \mathbf{follow}_m^f(p) \neq x). \quad (8.63)$$

Observe that for any natural  $n$ ,  $n = 0 \vee n > 0$  holds, which we can therefore add as a conjunct:

$$\iff \exists n. ((n = 0 \vee n > 0) \wedge \mathbf{follow}_n^f(p) = q \wedge \forall m < n. \mathbf{follow}_m^f(p) \neq x). \quad (8.64)$$

This simplifies as follows:

$$\iff \exists n. p = q \vee (n > 0 \wedge \mathbf{follow}_n^f(p) = q \wedge \forall m < n. \mathbf{follow}_m^f(p) \neq x) \quad (8.65)$$

$$\iff p = q \vee \exists n > 0. (\mathbf{follow}_n^f(p) = q \wedge \forall m < n. \mathbf{follow}_m^f(p) \neq x). \quad (8.66)$$

We replace  $n$  by  $n' + 1$  for natural  $n'$ :

$$\iff p = q \vee \exists n'. (\mathbf{follow}_{n'+1}^f(p) = q \wedge \forall m < n' + 1. \mathbf{follow}_m^f(p) \neq x). \quad (8.67)$$

As  $\mathbf{follow}_{n'+1}^f(p) = \mathbf{follow}_{n'}^f(p \rightarrow f)$ , this simplifies to

$$\iff p = q \vee \exists n'. (\mathbf{follow}_{n'}^f(p \rightarrow f) = q \wedge \forall m < n' + 1. \mathbf{follow}_m^f(p) \neq x). \quad (8.68)$$

By splitting the universal quantification into the two parts  $m = 0$  and  $m \geq 1$ , we obtain

$$\iff p = q \vee \exists n'. (\mathbf{follow}_{n'}^f(p \rightarrow f) = q \wedge p \neq x \wedge \forall 1 \leq m < n' + 1. \mathbf{follow}_m^f(p) \neq x). \quad (8.69)$$

The universal quantification is rewritten:

$$\iff p = q \vee \exists n'. (\mathbf{follow}_{n'}^f(p \rightarrow f) = q \wedge p \neq x \wedge \forall m < n'. \mathbf{follow}_m^f(p \rightarrow f) \neq x). \quad (8.70)$$

As the first and the third conjunct are equivalent to the definition of  $p \rightarrow f \xrightarrow[f]{x} q$ , the claim is shown.  $\blacksquare$

There are two simple consequences of rule A1:

$$p \xrightarrow[f]{x} p \quad \text{and} \quad p \xrightarrow[f]{p} q \iff p = q. \quad (8.71)$$

In the following example we use these consequences to prove (8.61), the reachability predicate formula for our first verification condition.

**Example 8.16.** Recall (8.61):

$$\text{list} \xrightarrow[0]{\text{next}} \text{list} \wedge \forall q \neq \text{list}. (\text{list} \xrightarrow[\text{list}]{\text{next}} q \implies q \rightarrow \text{payload} \neq a). \quad (8.72)$$

The first conjunct is a trivial instance of the first consequence. To show the second conjunct, we introduce a **Skolem variable**  $q'$  for the universal quantifier:<sup>9</sup>

$$(q' \neq \text{list} \wedge \text{list} \xrightarrow[\text{list}]{\text{next}} q') \implies q' \rightarrow \text{payload} \neq a. \quad (8.73)$$

By the second consequence, the left-hand side of the implication is FALSE. ▀

Even when the axioms are used, however, reasoning about a reachability predicate remains tedious. The goal is therefore to devise an automatic decision procedure for a logic that includes a reachability predicate. We mention several decision procedures for logics with reachability predicates in the bibliographical notes.

## 8.6 Problems

### 8.6.1 Pointer Formulas

**Problem 8.1 (semantics of pointer formulas).** Determine if the following pointer logic formulas are valid using the semantic translation:

1.  $x = y \implies \&x = \&y$ .
2.  $\&x \neq x$ .
3.  $\&x \neq \&y + i$ .
4.  $p \hookrightarrow x \implies *p = x$ .
5.  $p \hookrightarrow x \implies p \rightarrow f = x$ .
6.  $(p_1 \hookrightarrow p_2, x_1 \wedge p_2 \hookrightarrow \text{NULL}, x_2) \implies p_1 \neq p_2$ .

**Problem 8.2 (modeling dynamically allocated data structures).**

1. What data structure is modeled by  $\text{my-ds}(q, l)$  in the following? Draw an example.

$$\begin{aligned} \mathbf{c}(q, 0) &\doteq (*q).p = \text{NULL} \\ \mathbf{c}(q, i) &\doteq (*\text{list-elem}(q, i)).p = \text{list-elem}(q, i - 1) \quad \text{for } i \geq 1 \\ \mathbf{my-ds}(q, l) &\doteq \text{list-elem}(q, l) = \text{NULL} \wedge \forall 0 \leq i < l. \mathbf{c}(q, i) \end{aligned}$$

2. Write a recursive shorthand  $\text{DAG}(p)$  to denote that  $p$  points to the root of a directed acyclic graph.

<sup>9</sup> A Skolem variable is a ground variable introduced to eliminate a quantifier, i.e.,  $\forall x.P(x)$  is valid iff  $P(x')$  is valid for a new variable  $x'$ . This is a special case of Skolemization, which is named after Thoralf Skolem.

3. Write a recursive shorthand  $\mathbf{tree}(p)$  to denote that  $p$  points to the root of a tree.
4. Write a shorthand  $\mathbf{hashtbl}(p)$  to denote that  $p$  points to an array of lists.

**Problem 8.3 (extensions of the pointer logic).** Consider a pointer logic that only permits a conjunction of predicates of the following form, where  $p$  is a pointer, and  $f_i, g_i$  are field identifiers:

$$\forall p. p \rightarrow f_1 \rightarrow f_2 \rightarrow f_3 \dots = p \rightarrow g_1 \rightarrow g_2 \rightarrow g_3 \dots$$

Show that this logic is Turing complete.

**Problem 8.4 (axiomatization of the memory model).** Define a set of memory model axioms for an architecture that uses 32-bit integers and little-endian byte ordering.

**Problem 8.5 (partitioning the memory).** Suppose that a pointer logic permits restrictions on  $L$ , the memory layout. Give a counterexample to Theorem 8.14.

### 8.6.2 Reachability Predicates

**Problem 8.6 (semantics of reachability predicates).** Determine the satisfiability of the following reachability predicate formulas:

1.  $p \xrightarrow{f} q \wedge p \neq q$ .
2.  $p \xrightarrow{f} q \wedge p \xrightarrow{f} x$ .
3.  $p \xrightarrow{f} q \wedge q \xrightarrow{f} p$ .
4.  $\neg(p \xrightarrow{f} q) \wedge \neg(q \xrightarrow{f} p)$ .

**Problem 8.7 (modeling).** Try to write reachability predicate formulas for the following scenarios:

1.  $p$  points to a cyclic list where the next field is  $next$ .
2.  $p$  points to a NULL-terminated, doubly linked list.
3.  $p$  points to the root of a binary tree. The names of the fields for the left and right subtrees are  $l$  and  $r$ , respectively.
4.  $p$  points to the root of a binary tree as above, and the leaves are connected to a cyclic list.
5.  $p$  and  $q$  point to NULL-terminated singly linked lists that do not share cells.

**Problem 8.8 (decision procedures).** Build a decision procedure for a conjunction of atoms that have the form  $p \xrightarrow{f} q$  (or its negation).

**Problem 8.9 (program verification).** Write a code fragment that removes an element from a singly linked list, and provide the verification conditions using reachability predicate formulas.

## 8.7 Bibliographic Notes

The view of pointers as indices into a global array is commonplace, and similarly so is the identification of structure components with arrays. Leino's thesis is an instance of recent work applying this approach [117], and resembles our Sect. 8.3. An alternative point of view was proposed by Burstall: each component introduces an array, where the array indices are the addresses of the structures [42].

Transitive closure is frequently used to model recursive data structures. Immerman et al. explored the impact of adding transitive closure to a given logic. They showed that already very weak logics became undecidable as soon as transitive closure was added [101].

The PALE (Pointer Assertion Logic Engine) toolkit, implemented by Anders Møller, uses a graph representation for various dynamically allocated data structures. The graphs are translated into monadic second-order logic and passed to MONA, a decision procedure for this logic [129]. Michael Rabin proved in 1969 that the monadic second-order theory of trees was decidable [161].

The reachability predicate discussed in Sect. 8.5 was introduced by Greg Nelson [135]. This 1983 paper stated that the question of whether the set of (eight) axioms provided was complete remained open. A technical report gives a decision procedure for a conjunction of reachability predicates, which implies the existence of a complete axiomatization [138]. The procedure has linear time complexity.

Numerous modern logics are based on this idea. For example, Lahiri and Qadeer proposed two logics based on the idea of reachability predicates, and offered effective decision procedures [113, 114]. The decision procedure for [114] was based on a recent SMT solver.

Alain Deutsch [66] introduced an alias analysis algorithm that uses *symbolic access paths*, i.e., expressions that symbolically describe what field to follow for a given number of times. Symbolic access paths are therefore a generalization of the technique we described in Sect. 8.5. Symbolic access paths are very expressive when combined with an expressive logic for the *basis* of the access path, but this combination often results in undecidability.

Benedikt et al. [17] defined a logic for linked data structures. This logic uses constraints on paths (called *routing expressions*) in order to define memory

regions, and permits one to reason about sharing and reachability within such regions. These authors showed the logic to be decidable using a small-model property argument, but did not provide an efficient decision procedure.

A major technique for analyzing dynamically allocated data structures is *parametric shape analysis*, introduced by Sagiv, Reps, and Wilhelm [163, 173, 198]. An important concept in the shape analysis of Sagiv et al. is the use of Kleene's three-valued logic for distinguishing predicates that are true, false, or *unknown* in a particular abstract state. The resulting concretizations are more precise than an abstraction using traditional, two-valued logic.

*Separation Logic* (see the aside on this subject) was introduced by John Reynolds as an intuitionistic way of reasoning about dynamically allocated data structures [165]. Calcagno et al. [44] showed that deciding the validity of a formula in separation logic, even if robbed of its characteristic separating conjunction, was not recursively enumerable. On the other hand, they showed that once quantifiers were prohibited, validity became decidable. Decidable fragments of separation logic have been studied, for example by Berdine et al. [18, 19, 20]; these are typically restricted to predicates over lists. Parkinson and Bierman address the problem of modular reasoning about programs using separation logic [146].

Kuncak and Rinard introduced *regular graph constraints* as a representation of heaps. They showed that satisfiability of such heap summary graphs was decidable, whereas entailment was not [110].

Alias analysis techniques have also been integrated directly into verification algorithms. Manevich et al. described predicate abstraction techniques for singly linked lists [121]. Beyer et al. described how to combine a predicate abstraction tool that implements lazy abstraction with shape analysis [21]. Podelski and Wies propose *Boolean heaps* as an abstract model for heap-manipulating programs [157]. Here, the abstract domain is spanned by a vector of arbitrary first-order predicates characterizing the heap. Bingham and Rakamarić [24] also proposed to extend predicate abstraction with predicates designated to describe the heap. Distefano et al. [67] defined an abstract domain that is based on predicates drawn from separation logic. Berdine et al. use separation logic predicates in an add-on to Microsoft's SLAM device driver verifier, called TERMINATOR, in order to prove that loops iterating over dynamically allocated data structures terminated.

Most frameworks for reasoning about dynamically allocated memory treat the heap as composed of disjoint memory fragments, and do not model accesses beyond these fragments using pointer arithmetic. Calcagno et al. introduced a variant of separation logic that permits reasoning about low-level programs including pointer arithmetic [43]. This logic permits the analysis of infrastructure usually assumed to exist at higher abstraction layers, e.g., the code that implements the `malloc` function.



## 8.8 Glossary

The following symbols were used in this chapter:

<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$A$	Set of addresses	182
$D$	Set of data-words	182
$M$	Map from addresses to data-words	182
$L$	Memory layout	182
$\sigma(v)$	The size of $v$	182
$V$	Set of variables	182
$\llbracket \cdot \rrbracket$	Semantics of pointer expressions	187
$p \leftrightarrow z$	$p$ points to a variable with value $z$	187
$p \rightarrow f$	Shorthand for $(*p).f$	189
$\text{list}(p, l)$	$p$ points to a list of length $l$	190

---

## Quantified Formulas

### 9.1 Introduction

Quantification allows us to specify the extent of validity of a predicate, or in other words the **domain** (range of values) in which the predicate should hold. The syntactic element used in the logic for specifying quantification is called a **quantifier**. The most commonly used quantifiers are the *universal quantifier*, denoted by “ $\forall$ ”, and the *existential quantifier*, denoted by “ $\exists$ ”. These two quantifiers are interchangeable using the following equivalence:

$$\forall x. \varphi \iff \neg \exists x. \neg \varphi . \quad (9.1)$$

Some examples of quantified statements are:

- For any integer  $x$ , there is an integer  $y$  smaller than  $x$ :

$$\forall x \in \mathbb{Z}. \exists y \in \mathbb{Z}. y < x . \quad (9.2)$$

- There exists an integer  $y$  such that for any integer  $x$ ,  $x$  is greater than  $y$ :

$$\exists y \in \mathbb{Z}. \forall x \in \mathbb{Z}. x > y . \quad (9.3)$$

- (Bertrand’s Postulate) For any natural number greater than 1, there is a prime number  $p$  such that  $n < p < 2n$ :

$$\forall n \in \mathbb{N}. \exists p \in \mathbb{N}. n > 1 \implies (\text{isprime}(p) \wedge n < p < 2n) . \quad (9.4)$$

In these three examples, there is **quantifier alternation** between the universal and existential quantifiers. In fact, the satisfiability and validity problems that we considered in earlier chapters can be cast as decision problems for formulas with nonalternating quantifiers. When we ask whether the propositional formula

$$x \vee y \quad (9.5)$$

is satisfiable, we can equivalently ask whether *there exists* a truth assignment to  $x, y$  that satisfies this formula.<sup>1</sup> And when we ask whether

$$x > y \vee x < y \tag{9.6}$$

is valid for  $x, y \in \mathbb{N}$ , we can equivalently ask whether this formula holds for *all* naturals  $x$  and  $y$ . The formulations of these two decision problems, are, respectively,

$$\exists x \in \mathbb{B}. y \in \mathbb{B}. (x \vee y) \tag{9.7}$$

and

$$\forall x \in \mathbb{N}. \forall y \in \mathbb{N}. x > y \vee x < y. \tag{9.8}$$

We omit the domain of each quantified variable from now on when it is not essential for the discussion.

An important characteristic of quantifiers is the scope in which they are applied, called the **binding scope**. For example, in the following formula, the existential quantification over  $x$  overrides the external universal quantification over  $x$ :

$$\underbrace{\forall x. ((x < 0) \wedge \exists y. \overbrace{(y > x \wedge (y \geq 0 \vee \exists x. \underbrace{(y = x + 1)}_{\text{scope of } \exists x}))}_{\text{scope of } \exists y}))}_{\text{scope of } \forall x}}. \tag{9.9}$$

Within the scope of the second existential quantifier, all *occurrences* of  $x$  refer to the variable bound by the existential quantifier. It is impossible to refer directly to the variable bound by the universal quantifier. A possible solution is to rename  $x$  in the inner scope: clearly, this does not change the validity of the formula. After this renaming, we can assume that every occurrence of a variable is bound exactly once.

**Definition 9.1 (free variable).** *A variable is called free in a given formula if at least one of its occurrences is not bound by any quantifier.*

**Definition 9.2 (sentence).** *A formula  $\mathcal{Q}$  is called a sentence (or closed) if none of its variables is free.*

In this chapter we only focus on sentences.

Arbitrary first-order theories with quantifiers are undecidable. We restrict the discussion in this chapter to decidable theories only, and begin with two examples.

<sup>1</sup> As explained in Sect. 1.4.1, the difference between the two formulations, namely with no quantifiers and with nonalternating quantifiers, is that in the former all variables are free (unquantified), and hence a satisfying structure (a *model*) for such formulas includes an assignment to these variables. Since such assignments are necessary in many applications, this book uses the former formulation.

### 9.1.1 Example: Quantified Boolean Formulas

**Quantified propositional logic** is propositional logic enhanced with quantifiers. Sentences in quantified propositional logic are better known as **quantified Boolean formulas** (QBFs). The set of sentences permitted by the logic is defined by the following grammar:

$$\begin{aligned} \textit{formula} : & \textit{formula} \wedge \textit{formula} \mid \neg \textit{formula} \mid (\textit{formula}) \mid \\ & \textit{identifier} \mid \exists \textit{identifier}. \textit{formula} \end{aligned}$$

Other symbols such as “ $\vee$ ”, “ $\forall$ ” and “ $\iff$ ” can be constructed using elements of the formal grammar. Examples of quantified Boolean formulas are

- $\forall x. (x \vee \exists y. (y \vee \neg x))$ ,
- $\forall x. (\exists y. ((x \vee \neg y) \wedge (\neg x \vee y)) \wedge \exists z. ((\neg y \vee \neg x) \wedge (x \vee y)))$ .

### Complexity

The validity problem of QBF is PSPACE-complete, which means that it is theoretically harder to solve than SAT, which is “only” NP-complete<sup>2</sup>. Both of these problems (SAT and the QBF problem) are frequently presented as the quintessential problems of their respective complexity classes. The known algorithms for both problems are exponential.

### Usage example: chess

The following is an example of the use of QBF.

**Example 9.3.** QBF is a convenient way of modeling many finite two-player games. As an example, consider the problem of determining whether there is a winning strategy for a chess player in  $k$  steps, i.e., given a state of a board and assuming white goes first, can white take the black king in  $k$  steps, regardless of black’s moves? This problem can be modeled as QBF rather naturally, because what we ask is whether there *exists* a move of white such that *for all* possible moves of black that follow there *exists* a move of white such that *for all* possible moves of black... and so forth,  $k$  times, such that the goal of eliminating the black king is achieved. The number of steps  $k$  has to be an odd natural, as white plays both the first and last move.

<sup>2</sup> The difference between these two classes is that problems in NP are known to have nondeterministic algorithms that solve them in polynomial time. It has not been proven that these two classes are indeed different, but it is widely suspected that this is the case.

This is a classical problem in **planning**, a popular field of study in artificial intelligence. To formulate the chess problem in QBF<sup>3</sup>, we use the notation in Fig. 9.1. Every piece of each player has a unique index. Each location on the board has a unique index as well, and the location 0 of a piece indicates that it is outside the board. The size of the board is  $s$  (normally  $s = 8$ ), and hence there are  $s^2 + 1$  locations and  $4s$  pieces.

Symbol	Meaning
$x_{\{m,n,i\}}$	Piece $m$ is at location $n$ in step $i$ , for $1 \leq m \leq 4s$ , $0 \leq n \leq s^2$ , and $0 \leq i \leq k$ .
$I_0$	A set of clauses over the $x_{\{m,n,0\}}$ variables that represent the initial state of the board.
$T_i^w$	A set of clauses over the $x_{\{m,n,i\}}, x_{\{m,n,i+1\}}$ variables that represent the valid moves by white at step $i$ .
$T_i^b$	A set of clauses over the $x_{\{m,n,i\}}, x_{\{m,n,i+1\}}$ variables that represent the valid moves by black at step $i$ .
$G_k$	A set of clauses over the $x_{\{m,n,k\}}$ variables that represent the goal, i.e., in step $k$ the black king is off the board and the white king is on the board.

**Fig. 9.1.** Notation used in Example 9.3

We use the following convention: we write  $\{x_{m,n,i}\}$  to represent the set of variables  $\{x_{\{m,n,i\}} \mid m, n, i \text{ in their respective ranges}\}$ . Let us begin with the following attempt to formulate the problem:

$$\begin{aligned} & \exists\{x_{\{m,n,0\}}\} \exists\{x_{\{m,n,1\}}\} \forall\{x_{\{m,n,2\}}\} \exists\{x_{\{m,n,3\}}\} \cdots \forall\{x_{\{m,n,k-1\}}\} \exists\{x_{\{m,n,k\}}\} \cdot \\ & I_0 \wedge (T_0^w \wedge T_2^w \wedge \cdots \wedge T_{k-1}^w) \wedge (T_1^b \wedge T_3^b \wedge \cdots \wedge T_{k-2}^b) \wedge G_k . \end{aligned} \tag{9.10}$$

This formulation includes the necessary restrictions on the initial and goal states, as well as on the allowed transitions. The problem is that this formula is not valid for any initial configuration, because black can make an illegal move – such as moving two pieces at once – which falsifies the formula (it contradicts the subformula  $T_i$  for some odd  $i$ ).

The formula needs to be weakened, as it is sufficient to find a white move for the *legal* moves of black:

<sup>3</sup> Classical formulations of planning problems distinguish between actions (moves in this case) and states. Here we chose to present a formulation based on states only.

$$\exists\{x_{\{m,n,0\}}\}\exists\{x_{\{m,n,1\}}\}\forall\{x_{\{m,n,2\}}\}\exists\{x_{\{m,n,3\}}\}\cdots\forall\{x_{\{m,n,k-1\}}\}\exists\{x_{\{m,n,k\}}\}. \\ I_0 \wedge ((T_1^b \wedge T_3^b \wedge \cdots \wedge T_{k-2}^b) \implies (T_0^w \wedge T_2^w \wedge \cdots \wedge T_{k-1}^w \wedge G_k)). \quad (9.11)$$

Is this formula a faithful representation of the chess problem? Unfortunately not, because of the possibility of a stalemate: there could be a situation in which black is not able to make a valid move, which results in a remis. A possible solution for this problem is to ban white from making moves that result in such a state by modifying  $T^w$  appropriately. ■

### 9.1.2 Example: Quantified Disjunctive Linear Arithmetic

The syntax of **quantified disjunctive linear arithmetic** (QDLA) is defined by the following grammar:

$$\begin{aligned} \text{formula} &: \text{formula} \wedge \text{formula} \mid \neg \text{formula} \mid (\text{formula}) \mid \\ &\quad \text{predicate} \mid \forall \text{identifier. formula} \\ \text{predicate} &: \Sigma_i a_i x_i \leq c \end{aligned}$$

where  $c$  and  $a_i$  are constants, for all  $i$ . The domain of the variables (identifiers) is the reals. As before, other symbols such as “ $\vee$ ”, “ $\exists$ ” and “ $=$ ” can be defined using the formal grammar.

#### Aside: Presburger Arithmetic

Presburger arithmetic has the same grammar as quantified disjunctive linear arithmetic, but is defined over the natural numbers rather than over the reals. Presburger arithmetic is decidable and, as proven by Fischer and Rabin [75], there is a lower bound of  $2^{2^{c \cdot n}}$  on the worst-case run-time complexity of a decision procedure for this theory, where  $n$  is the length of the input formula and  $c$  is a positive constant. This theory is named after Mojzesz Presburger, who introduced it in 1929 and proved its decidability. Replacing the Fourier–Motzkin procedure with the Omega test (see Sect. 5.5) in the procedure described in this section gives a decision procedure for this theory. Other decision procedures for Presburger arithmetic are mentioned in the bibliographic notes at the end of this chapter.

As an example, the following is a QDLA formula:

$$\forall x. \exists y. \exists z. (y + 1 \leq x \quad \vee \quad z + 1 \leq y \quad \wedge \quad 2x + 1 \leq z). \quad (9.12)$$

## 9.2 Quantifier Elimination

### 9.2.1 Prenex Normal Form

We begin by defining a normal form for quantified formulas.

**Definition 9.4 (prenex normal form).** A formula is said to be in prenex normal form (PNF) if it is in the form

$$Q[n]V[n] \cdots Q[1]V[1]. \langle \text{quantifier-free formula} \rangle, \quad (9.13)$$

where for all  $i \in \{1, \dots, n\}$ ,  $Q[i] \in \{\forall, \exists\}$  and  $V[i]$  is a variable.

We call the quantification string on the left of the formula the **quantification prefix**, and call the quantifier-free formula to the right of the quantification prefix the **quantification suffix** (also called the *matrix*).

**Lemma 9.5.** For every quantified formula  $\mathcal{Q}$  there exists a formula  $\mathcal{Q}'$  in prenex normal form such that  $\mathcal{Q}$  is valid if and only if  $\mathcal{Q}'$  is valid.

Algorithm 9.2.1 transforms an input formula into prenex normal form.

**Algorithm 9.2.1: PRENEX**

**Input:** A quantified formula

**Output:** A formula in prenex normal form

1. Eliminate Boolean connectives other than  $\vee, \wedge, \neg$ .
2. Push negations to the right across all quantifiers, using De Morgan's rules (see Sect. 1.3) and (9.1).
3. If there are name conflicts across scopes, solve by renaming: give each variable in each scope a unique name.
4. Move quantifiers out by using equivalences such as

$$\begin{aligned} \phi_1 \wedge Qx. \phi_2(x) &\iff Qx. (\phi_1 \wedge \phi_2(x)), \\ \phi_1 \vee Qx. \phi_2(x) &\iff Qx. (\phi_1 \vee \phi_2(x)), \\ Q_1y. \phi_1(y) \wedge Q_2x. \phi_2(x) &\iff Q_1y. Q_2x. (\phi_1(y) \wedge \phi_2(x)), \\ Q_1y. \phi_1(y) \vee Q_2x. \phi_2(x) &\iff Q_1y. Q_2x. (\phi_1(y) \vee \phi_2(x)), \end{aligned}$$

where  $Q, Q_1, Q_2 \in \{\forall, \exists\}$  are quantifiers,  $x \notin \text{vars}(\phi_1)$ , and  $y \notin \text{vars}(\phi_2)$ .

**Example 9.6.** We demonstrate Algorithm 9.2.1 with the following formula:

$$\mathcal{Q} := \neg \exists x. \neg (\exists y. ((y \implies x) \wedge (\neg x \vee y)) \wedge \neg \forall y. ((y \wedge x) \vee (\neg x \wedge \neg y))). \quad (9.14)$$

In steps 1 and 2, eliminate “ $\implies$ ” and push negations inside:

$$\forall x. (\exists y. ((\neg y \vee x) \wedge (\neg x \vee y)) \wedge \exists y. ((\neg y \vee \neg x) \wedge (x \vee y))). \quad (9.15)$$

In step 3, rename  $y$  as there are two quantifications over this variable:

$$\forall x. (\exists y_1. ((\neg y_1 \vee x) \wedge (\neg x \vee y_1)) \wedge \exists y_2. ((\neg y_2 \vee \neg x) \wedge (x \vee y_2))) . \quad (9.16)$$

Finally, in step 4, move quantifiers to the left of the formula:

$$\forall x. \exists y_1. \exists y_2. (\neg y_1 \vee x) \wedge (\neg x \vee y_1) \wedge (\neg y_2 \vee \neg x) \wedge (x \vee y_2) . \quad (9.17)$$

▀

We assume from here on that the input formula is given in prenex normal form.

### 9.2.2 Quantifier Elimination Algorithms

A *quantifier elimination* algorithm transforms a quantified formula into an equivalent formula without quantifiers.<sup>4</sup> The procedures that we present next require that all the quantifiers are eliminated in order to check for validity.

It is sufficient to show that there exists a procedure for eliminating an existential quantifier. Universal quantifiers can be eliminated by making use of (9.1). For this purpose we define a general notion of *projection*, which has to be concretized for each individual theory.

**Definition 9.7 (projection).** A projection of a variable  $x$  from a quantified formula in prenex normal form with  $n$  quantifiers,

n

$$\mathcal{Q}_1 = Q[n]V[n] \cdots Q[2]V[2]. \exists x. \phi , \quad (9.18)$$

is a formula

$$\mathcal{Q}_2 = Q[n]V[n] \cdots Q[2]V[2]. \phi' \quad (9.19)$$

(where both  $\phi$  and  $\phi'$  are quantifier-free), such that  $x \notin \text{var}(\phi')$ , and  $\mathcal{Q}_1$  and  $\mathcal{Q}_2$  are logically equivalent.

Given a projection algorithm PROJECT, Algorithm 9.2.2 eliminates all quantifiers. Assuming that we begin with a sentence (see Definition 9.2), the remaining formula is over constants and easily solvable.

<sup>4</sup> Every sentence is equivalent to a formula without quantifiers, namely TRUE or FALSE. But this does not mean that every theory has a quantifier elimination algorithm. The existence of a quantifier elimination algorithm typically implies the decidability of the logic.



**Algorithm 9.2.2: QUANTIFIER ELIMINATION**

**Input:** A sentence  $Q[n]V[n] \cdots Q[1]V[1]. \phi$ , where  $\phi$  is quantifier-free

**Output:** A (quantifier-free) formula over constants  $\phi'$ , which is valid if and only if  $\phi$  is valid

1.  $\phi' := \phi$ ;
2. **for**  $i := 1, \dots, n$  **do**
3.   **if**  $Q[i] = \exists$  **then**
4.      $\phi' := \text{PROJECT}(\phi', V[i])$ ;
5.   **else**
6.      $\phi' := \neg\text{PROJECT}(\neg\phi', V[i])$ ;
7. **Return**  $\phi'$ ;

We now show two examples of projection procedures and their use in quantifier elimination.

**9.2.3 Quantifier Elimination for Quantified Boolean Formulas**

Eliminating an existential quantifier over a conjunction of Boolean literals is trivial: if  $x$  appears with both phases in the conjunction, then the formula is unsatisfiable; otherwise,  $x$  can be removed. For example,

$$\begin{aligned} \exists y. \exists x. x \wedge \neg x \wedge y &= \text{FALSE}, \\ \exists y. \exists x. x \wedge y &= \exists y. y = \text{TRUE}. \end{aligned} \tag{9.20}$$

This observation can be used if we first convert the quantification suffix to DNF and then apply projection to each term separately. This is justified by the following equivalence:

$$\exists x. \bigvee_i \bigwedge_j l_{ij} \iff \bigvee_i \exists x. \bigwedge_j l_{ij}, \tag{9.21}$$

where  $l_{ij}$  are literals. But since converting formulas to DNF can result in an exponential growth in the formula size (see Sect. 1.16), it is preferable to have a projection that works directly on the CNF, or better yet, on a general Boolean formula. We consider two techniques: *binary resolution* (see Definition 2.11), which works directly on CNF formulas, and *expansion*.

**Projection with Binary Resolution**

Resolution gives us a method to eliminate a variable  $x$  from a pair of clauses in which  $x$  appears with opposite phases. To eliminate  $x$  from a CNF formula by projection (Definition 9.7), we need to apply resolution to *all* pairs of clauses

where  $x$  appears with opposite phases. This eliminates  $x$  together with its quantifier. For example, given the formula

$$\exists y. \exists z. \exists x. (y \vee x) \wedge (z \vee \neg x) \wedge (y \vee \neg z \vee \neg x) \wedge (\neg y \vee z), \quad (9.22)$$

we can eliminate  $x$  together with  $\exists x$  by applying resolution on  $x$  to the first and second clauses, and to the first and third clauses, resulting in:

$$\exists y. \exists z. (y \vee z) \wedge (y \vee \neg z) \wedge (\neg y \vee z). \quad (9.23)$$

What about universal quantifiers? Relying on (9.1), in the case of CNF formulas, results in a surprisingly easy shortcut to eliminating universal quantifiers: simply erase them from the formula. For example, eliminating  $x$  and  $\forall x$  from

$$\exists y. \exists z. \forall x. (y \vee x) \wedge (z \vee \neg x) \wedge (y \vee \neg z \vee \neg x) \wedge (\neg y \vee z) \quad (9.24)$$

results in

$$\exists y. \exists z. (y) \wedge (z) \wedge (y \vee \neg z) \wedge (\neg y \vee z). \quad (9.25)$$

This step is called **forall reduction**. It should be applied only after removing tautology clauses (clauses in which a literal appears with both phases). We leave the proof of correctness of this trick to Problem 9.3. Intuitively, however, it is easy to see why this is correct: if the formula is evaluated to TRUE for all values of  $x$ , this means that we cannot satisfy a clause while relying on a specific value of  $x$ .

**Example 9.8.** In this example, we show how to use resolution on both universal and existential quantifiers. Consider the following formula:

$$\begin{aligned} & \forall u_1. \forall u_2. \exists e_1. \forall u_3. \exists e_3. \exists e_2. \\ & (u_1 \vee \neg e_1) \wedge (\neg u_1 \vee \neg e_2 \vee e_3) \wedge (u_2 \vee \neg u_3 \vee \neg e_1) \wedge (e_1 \vee e_2) \wedge (e_1 \vee \neg e_3). \end{aligned} \quad (9.26)$$

By resolving the second and fourth clauses on  $e_2$ , we obtain

$$\begin{aligned} & \forall u_1. \forall u_2. \exists e_1. \forall u_3. \exists e_3. \\ & (u_1 \vee \neg e_1) \wedge (\neg u_1 \vee e_1 \vee e_3) \wedge (u_2 \vee \neg u_3 \vee \neg e_1) \wedge (e_1 \vee \neg e_3). \end{aligned} \quad (9.27)$$

By resolving the second and fourth clauses on  $e_3$ , we obtain

$$\forall u_1. \forall u_2. \exists e_1. \forall u_3. (u_1 \vee \neg e_1) \wedge (\neg u_1 \vee e_1) \wedge (u_2 \vee \neg u_3 \vee \neg e_1). \quad (9.28)$$

By eliminating  $u_3$ , we obtain

$$\forall u_1. \forall u_2. \exists e_1. (u_1 \vee \neg e_1) \wedge (\neg u_1 \vee e_1) \wedge (u_2 \vee \neg e_1). \quad (9.29)$$

By resolving the first and second clauses on  $e_1$ , and the second and third clauses on  $e_1$ , we obtain

$$\forall u_1. \forall u_2. (u_1 \vee \neg u_1) \wedge (\neg u_1 \vee u_2). \quad (9.30)$$

The first clause is a tautology and hence is removed. Next,  $u_1$  and  $u_2$  are removed, which leaves us with the empty clause. The formula, therefore, is not valid.  $\blacksquare$

What is the complexity of this procedure? Consider the elimination of a quantifier  $\exists x$ . In the worst case, half of the clauses contain  $x$  and half  $\neg x$ . Since we create a new clause from each pair of the two types of clauses, this results in  $O(m^2)$  new clauses, while we erase the  $m$  old clauses that contain  $x$ . Repeating this process  $n$  times results in  $O(m^{2^n})$  clauses.

This seems to imply that the complexity of projection with binary resolution is doubly exponential. This, in fact, is only true if we do not prevent duplicate clauses. Observe that there cannot be more than  $3^N$  distinct clauses, where  $N$  is the total number of variables. The reason is that each variable can appear positively, negatively, or not at all in a clause. This implies that if we add each clause at most once, the number of clauses is only singly-exponential in  $n$  (assuming  $N$  is not exponentially larger than  $n$ ).

$N$

### Expansion-Based Quantifier Elimination

The following quantifier elimination technique is based on expansion of quantifiers, according to the following equivalences:

$$\exists x. \varphi = \varphi|_{x=0} \vee \varphi|_{x=1} , \quad (9.31)$$

$$\forall x. \varphi = \varphi|_{x=0} \wedge \varphi|_{x=1} . \quad (9.32)$$

The notation  $\varphi|_{x=0}$  (the *restrict* operation; see p. 46) simply means that  $x$  is replaced with 0 (FALSE) in  $\varphi$ . Note that (9.32) can be derived from (9.31) by using (9.1).

Projections using expansion result in formulas that grow to  $O(m \cdot 2^n)$  clauses in the worst case, where, as before,  $m$  is the number of clauses in the original formula. In contrast to binary resolution, there is no need to refrain from using duplicate clauses in order to remain singly exponential in  $n$ . Furthermore, this technique can be applied directly to non-CNF formulas, in contrast to resolution, as the following example shows.

**Example 9.9.** Consider the following formula:

$$\exists y. \forall z. \exists x. (y \vee (x \wedge z)) . \quad (9.33)$$

Applying (9.31) to  $\exists x$  results in

$$\exists y. \forall z. (y \vee (x \wedge z))|_{x=0} \vee (y \vee (x \wedge z))|_{x=1} , \quad (9.34)$$

which simplifies to

$$\exists y. \forall z. (y \vee z) . \quad (9.35)$$

Applying (9.32) yields

$$\exists y. (y \vee z)|_{z=0} \wedge (y \vee z)|_{z=1} , \quad (9.36)$$

which simplifies to

$$\exists y. (y) , \quad (9.37)$$

which is obviously valid. Hence, (9.33) is valid.  $\blacksquare$

### 9.2.4 Quantifier Elimination for Quantified Disjunctive Linear Arithmetic

Once again we need a projection method. We use the Fourier–Motzkin elimination, which was described in Sect. 5.4. This technique resembles the resolution method introduced in Sect. 9.2.3, and has a worst-case complexity of  $O(m^{2^n})$ . It can be applied directly to a conjunction of linear atoms and, consequently, if the input formula has an arbitrary structure, it has to be converted first to DNF.

Let us briefly recall the Fourier–Motzkin elimination method. In order to eliminate a variable  $x_n$  from a formula with variables  $x_1, \dots, x_n$ , for every two conjoined constraints of the form

$$\sum_{i=1}^{n-1} a'_i \cdot x_i < x_n < \sum_{i=1}^{n-1} a_i \cdot x_i, \quad (9.38)$$

where for  $i \in \{1, \dots, n-1\}$ ,  $a_i$  and  $a'_i$  are constants, we generate a new constraint

$$\sum_{i=1}^{n-1} a'_i \cdot x_i < \sum_{i=1}^{n-1} a_i \cdot x_i. \quad (9.39)$$

After generating all such constraints for  $x_n$ , we remove all constraints that involve  $x_n$  from the formula.

**Example 9.10.** Consider the following formula:

$$\forall x. \exists y. \exists z. (y + 1 \leq x \quad \wedge \quad z + 1 \leq y \quad \wedge \quad 2x + 1 \leq z). \quad (9.40)$$

By eliminating  $z$ , we obtain

$$\forall x. \exists y. (y + 1 \leq x \quad \wedge \quad 2x + 1 \leq y - 1). \quad (9.41)$$

By eliminating  $y$ , we obtain

$$\forall x. (2x + 2 \leq x - 1). \quad (9.42)$$

Using (9.1), we obtain

$$\neg \exists x. \neg(2x + 2 \leq x - 1), \quad (9.43)$$

or, equivalently,

$$\neg \exists x. x > -3, \quad (9.44)$$

which is obviously not valid. ▀

### 9.3 Search-Based Algorithms for Quantified Boolean Formulas

Most competitive QBF solvers are based on an adaptation of DPLL solvers. The adaptation that we consider here is naive, in that it resembles the basic DPLL algorithm without the more advanced features such as learning and nonchronological backtracking (see Chap. 2 for details of DPLL solvers).

The key difference between SAT and the QBF problem is that the latter requires handling of quantifier alternation. The binary search tree now has to distinguish between **universal nodes** and **existential nodes**. Universal nodes are labeled with a symbol “ $\forall$ ”, as can be seen in the right-hand drawing in Fig. 9.2.



**Fig. 9.2.** An existential node (*left*) and a universal node (*right*) in a QBF search-tree

A QBF binary search tree corresponding to a QBF  $\mathcal{Q}$ , is defined as follows.

**Definition 9.11 (QBF search tree corresponding to a quantified Boolean formula).** *Given a QBF  $\mathcal{Q}$  in prenex normal form and an ordering of its variables (say,  $x_1, \dots, x_n$ ), a QBF search tree corresponding to  $\mathcal{Q}$  is a binary labeled tree of height  $n + 1$  with two types of internal nodes, universal and existential, in which:*

- *The root node is labeled with  $\mathcal{Q}$  and associated with depth 0.*
- *One of the children of each node at level  $i$ ,  $0 \leq i < n$ , is marked with  $x_{i+1}$ , and the other with  $\neg x_{i+1}$ .*
- *A node in level  $i$ ,  $0 \leq i < n$ , is universal if the variable in level  $i + 1$  is universally quantified.*
- *A node in level  $i$ ,  $0 \leq i < n$ , is existential if the variable in level  $i + 1$  is existentially quantified.*

The validity of a QBF tree is defined recursively, as follows.


**Definition 9.12 (validity of a QBF tree).** *A QBF tree is valid if its root is satisfied. This is determined recursively according to the following rules:*

- *A leaf in a QBF binary tree corresponding to a QBF  $\mathcal{Q}$  is satisfied if the assignment corresponding to the path to this leaf satisfies the quantification suffix of  $\mathcal{Q}$ .*
- *A universal node is satisfied if both of its children are satisfied.*
- *An existential node is satisfied if at least one of its children is satisfied.*

**Example 9.13.** Consider the formula

$$\mathcal{Q} := \exists e. \forall u. (e \vee u) \wedge (\neg e \vee \neg u). \quad (9.45)$$

The corresponding QBF tree appears in Fig. 9.3.



**Fig. 9.3.** A QBF search tree for the formula  $\mathcal{Q}$  of (9.45)

The second and third  $u$  nodes are the only nodes that are satisfied (since  $(e, \neg u)$  and  $(\neg e, u)$  are the only assignments that satisfy the suffix). Their parent nodes,  $e$  and  $\neg e$ , are not satisfied, because they are universal nodes and only one of their child nodes is satisfied. In particular, the root node, representing  $\mathcal{Q}$ , is not satisfied and hence  $\mathcal{Q}$  is not valid.  $\blacksquare$

A naive implementation based on these ideas is described in Algorithm 9.3.1. More sophisticated algorithms exist [208, 209], in which techniques such as nonchronological backtracking and learning are applied: as in SAT, in the QBF problem we are also not interested in searching the whole search space defined by the above graph, but rather in pruning it as much as possible.

The notation  $\phi|_{\hat{v}}$  in line 6 refers to the simplification of  $\phi$  resulting from the assignments in the assignment set  $\hat{v}$ .<sup>5</sup> For example, let  $\hat{v} := \{x \mapsto 0, y \mapsto 1\}$ . Then

$$(x \vee (y \wedge z))|_{\hat{v}} = (z). \quad (9.46)$$

**Example 9.14.** Consider (9.45) once again:

$$\mathcal{Q} := \exists e. \forall u. (e \vee u) \wedge (\neg e \vee \neg u).$$

The progress of Algorithm 9.3.1 when applied to this formula, with the variable ordering  $u, e$ , is shown in Fig. 9.4.  $\blacksquare$

<sup>5</sup> This notation represents an extension of the *restrict* operation that was introduced on p. 46, from an assignment of a single variable to an assignment of a set of variables.

**Algorithm 9.3.1: SEARCH-BASED DECISION OF QBF**

**Input:** A QBF  $\mathcal{Q}$  in PNF  $Q[n]V[n] \cdots Q[1]V[1]$ .  $\phi$ , where  $\phi$  is in CNF

**Output:** “Valid” if  $\mathcal{Q}$  is valid, and “Not valid” otherwise

```

1. function MAIN(QBF formula  $\mathcal{Q}$ )
2.   if QBF( $\mathcal{Q}, \emptyset, n$ ) then return “Valid”;
3.   else return “Not valid”;
4.
5. function BOOLEAN QBF( $\mathcal{Q}$ , assignment set  $\hat{v}$ , int level)
6.   if ( $\phi|_{\hat{v}}$  simplifies to FALSE) then return FALSE;
7.   if (level = 0) then return TRUE;
8.   if ( $Q[\text{level}] = \forall$ ) then
9.     return  $\left( \text{QBF}(\mathcal{Q}, \hat{v} \cup \neg V[\text{level}], \text{level} - 1) \wedge \right.$ 
10.     $\left. \text{QBF}(\mathcal{Q}, \hat{v} \cup V[\text{level}], \text{level} - 1) \right)$ ;
11.  else
12.    return  $\left( \text{QBF}(\mathcal{Q}, \hat{v} \cup \neg V[\text{level}], \text{level} - 1) \vee \right.$ 
13.     $\left. \text{QBF}(\mathcal{Q}, \hat{v} \cup V[\text{level}], \text{level} - 1) \right)$ ;

```

## 9.4 Problems

### 9.4.1 Warm-up Exercises

**Problem 9.1 (example of forall reduction).** Show that the equivalence

$$\exists e. \exists f. \forall u. (e \vee f \vee u) \equiv \exists e. \exists f. (e \vee f) \quad (9.47)$$

holds.

**Problem 9.2 (expansion-based quantifier elimination).** Is the following formula valid? Check by eliminating all quantifiers with expansion. Perform simplifications when possible.

$$\mathcal{Q} := \forall x_1. \forall x_2. \forall x_3. \exists x_4. (x_1 \implies (x_2 \implies x_3)) \implies ((x_1 \wedge x_2 \implies x_3) \wedge (x_4 \vee x_1)). \quad (9.48)$$

### 9.4.2 QBF

**Problem 9.3 (eliminating universal quantifiers from CNF).** Let

$$\mathcal{Q} := Q[n]V[n] \cdots Q[2]V[2]. \forall x. \phi, \quad (9.49)$$

where  $\phi$  is a CNF formula. Let

Recursion level	Line	Comment
0	2	$QBF(\mathcal{Q}, \emptyset, 2)$ is called.
0	6,7	The conditions in these lines do not hold.
0	8	$Q[2] = \exists$ .
0	11	$QBF(\mathcal{Q}, \{e = 0\}, 1)$ is called first.
1	6	$\phi _{e=0} = (u)$ .
1	8	$Q[1] = \forall$ .
1	9	$QBF(\mathcal{Q}, \{e = 0, u = 0\}, 0)$ is called first.
2	6	$\phi _{e=0, u=0} = \text{FALSE}$ . return FALSE.
1	9	return FALSE.
0	11	$QBF(\mathcal{Q}, \{e = 1\}, 1)$ is called second.
1	6	$\phi _{e=1} = (\neg u)$ .
1	8	$Q[1] = \forall$ .
1	9	$QBF(\mathcal{Q}, \{e = 1, u = 0\}, 0)$ is called first.
2	6	$\phi _{e=1, u=0} = \text{TRUE}$ .
2	7	return TRUE.
1	9	$QBF(\mathcal{Q}, \{e = 1, u = 1\}, 0)$ is called second.
2	6	$\phi _{e=1, u=1} = \text{FALSE}$ ; return FALSE.
1	9	return FALSE.
0	11	return FALSE.
0	3	return "Not valid".

**Fig. 9.4.** A trace of Algorithm 9.3.1 when applied to (9.45)

$$\mathcal{Q}' := Q[n]V[n] \cdots Q[2]V[2]. \phi', \quad (9.50)$$

where  $\phi'$  is the same as  $\phi$  except that  $x$  and  $\neg x$  are erased from all clauses.

1. Prove that  $\mathcal{Q}$  and  $\mathcal{Q}'$  are logically equivalent if  $\phi$  does not contain tautology clauses.
2. Show an example where  $\mathcal{Q}$  and  $\mathcal{Q}'$  are not logically equivalent if  $\phi$  contains tautology clauses.

**Problem 9.4 (modeling: the diameter problem).** QBFs can be used for finding the longest shortest path of any state from an initial state in a finite state machine. More formally, what we would like to find is defined as follows:

**Definition 9.15 (initialized diameter of a finite state machine).** *The initialized diameter of a state machine is the smallest  $k \in \mathbb{N}$  for which every node reachable in  $k + 1$  steps can also be reached in  $k$  steps or fewer.*

Our assumption is that the finite state machine is too large to represent or explore explicitly: instead, it is given to us implicitly in the form of a transition system, in a similar fashion to the chess problem that was described in Sect. 9.1.1.

For the purpose of this problem, a finite transition system is a tuple  $\langle S, I, T \rangle$ , where  $S$  is a finite set of states, each of which is a valuation of



a finite set of variables ( $V \cup V' \cup In$ ).  $V$  is the set of state variables and  $V'$  is the corresponding set of next-state variables.  $In$  is the set of input variables.  $I$  is a predicate over  $V$  defining the initial states, and  $T$  is a transition function that maps each variable  $v \in V'$  to a predicate over  $V \cup I$ .

An example of a class of state machines that are typically represented in this manner is digital circuits. The initialized diameter of a circuit is important in the context of formal verification: it represents the largest depth to which one needs to search for an error state.

Given a transition system  $M$  and a natural  $k$ , formulate with QBF the problem of whether  $k$  is the diameter of the graph represented by  $M$ . Introduce proper notation in the style of the chess problem that was described in Sect. 9.1.1.

**Problem 9.5 (search-based QBFs).** Apply Algorithm 9.3.1 to the formula

$$\mathcal{Q} := \forall u. \exists e. (e \vee u)(\neg e \vee \neg u). \quad (9.51)$$

Show a trace of the algorithm as in Fig. 9.4.

**Problem 9.6 (QBFs and resolution).** Using resolution, check whether the formula

$$\mathcal{Q} := \forall u. \exists e. (e \vee u)(\neg e \vee \neg u) \quad (9.52)$$

is valid.

**Problem 9.7 (projection by resolution).** Show that the pairwise resolution suggested in Sect. 9.2.3 results in a projection as defined in Definition 9.7.

**Problem 9.8 (QBF refutations).** Let

$$\mathcal{Q} = Q[n]V[n]. \cdots Q[1]V[1]. \phi, \quad (9.53)$$

where  $\phi$  is in CNF and  $\mathcal{Q}$  is FALSE, i.e.,  $\mathcal{Q}$  is not valid. Propose a *proof format* for such QBFs that is generally applicable, i.e., allows us to give a proof for any QBF that is not valid (similarly to the way that binary-resolution proofs provide a proof format for propositional logic).

**Problem 9.9 (QBF models).** Let

$$\mathcal{Q} = Q[n]V[n]. \cdots Q[1]V[1]. \phi, \quad (9.54)$$

where  $\phi$  is in CNF and  $\mathcal{Q}$  is TRUE, i.e.,  $\mathcal{Q}$  is valid. In contrast to the quantifier-free SAT problem, we cannot provide a satisfying assignment to all variables that convinces us of the validity of  $\mathcal{Q}$ .

- Propose a *proof format* for valid QBFs.
- Provide a proof for the formula in Problem 9.6 using your proof format.
- Provide a proof for the following formula:

$$\forall u. \exists e. (u \vee \neg e)(\neg u \vee e).$$

## 9.5 Bibliographic Notes

Stockmeyer and his PhD advisor at MIT, Meyer, identified the QBF problem as PSPACE-complete as part of their work on the polynomial hierarchy [184, 185]. The idea of solving QBF by alternating between resolution and eliminating universally quantified variables from CNF clauses was proposed by Büning, Karpinski and Flögel [41]. The resolution part was termed **Q-resolution** (recall that the original SAT-solving technique developed by Davis and Putnam was based on resolution [57]).

There are many similarities in the research directions of SAT and QBF, and in fact there are researchers who are active in both areas. The positive impact that annual competitions and benchmark repositories have had on the development of SAT solvers, has led to similar initiatives for the QBF problem (e.g., see QBFLIB [85], which at the beginning of 2008 included more than 13 000 examples and a collection of more than 50 QBF solvers). Further, similarly to the evidence provided by propositional SAT solvers (namely a satisfying assignment or a resolution proof), many QBF solvers now provide a **certificate** of the validity or invalidity of a QBF instance [103] (also see Problems 9.8 and 9.9). Not surprisingly, there is a huge difference between the size of problems that can be solved in a reasonable amount of time by the best QBF solvers (thousands or a few tens of thousands of variables) and the size of problems that can be solved by the best SAT solvers (several hundreds of thousands or even a few millions of variables). It turns out that the exact encoding of a given problem can have a very significant impact on the ability to solve it – see, for example, the work by Sabharwal et al. [172]. The formulation of the chess problem in this chapter is inspired by that paper.

The research in the direction of applying propositional SAT techniques to QBFs, such as adding conflict and blocking clauses and the search-based method, is mainly due to work by Zhang and Malik [208, 209]. Quantifier expansion is folk knowledge, and was used for efficient QBF solving by, for example, Biere [22]. A similar type of expansion, called Shannon expansion, was used for one-alternation QBFs in the context of symbolic *model checking* with BDDs – see, for example, the work of McMillan [125]. Variants of BDDs were used for QBF solving in [83].

Presburger arithmetic is due to Mojżesz Presburger, who published his work, in German, in 1929 [159]. At that time, Hilbert considered Presburger’s decidability result as a major step towards full mechanization of mathematics (full mechanization of mathematics was the ultimate goal of many mathematicians, such as Leibniz and Peano, much earlier than that), which later on proved to be an impossibility, owing to Gödel’s incompleteness theorem. Gödel’s result refers to **Peano arithmetic**, which is the same as Presburger arithmetic with the addition of multiplication. One of the first mechanical deduction systems was an implementation of Presburger’s algorithm on the Johnniac, a vacuum tube computer, in 1954. At the time, it was considered

a major step that the program was able to show that the sum of two even numbers is an even number.

Two well-known approaches for solving Presburger formulas, in addition to the one based on the Omega test that was mentioned in this chapter, are due to Cooper [51] and the family of methods based on finite automata and model checking: see the article by Wolper and Boigelot [203] and the publications regarding the LASH system, as well as Ganesh, Berezin, and Dill's survey and empirical comparison [77] of such methods when applied to unquantified Presburger formulas.

The problem of deciding quantified formulas over nonlinear real arithmetic is decidable, although a description of a decision procedure for this problem is not within the scope of this book. A well-known decision procedure for this theory is cylindrical algebraic decomposition (CAD). A comparison of CAD with other techniques can be found in [68]. Several tutorials on CAD can be found on the Web.

## 9.6 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$\forall, \exists$	The universal and existential quantification symbols	207
$n$	The number of quantifiers	213
$N$	The total number of variables (not only those existentially quantified)	216
$\phi _{\hat{v}}$	A simplification of $\phi$ based on the assignments in $\hat{v}$ . This extends the <i>restrict</i> operator (p. 46)	219

---

## Deciding a Combination of Theories

### 10.1 Introduction

The decision procedures that we have studied so far focus on one specific theory. Verification conditions that arise in practice, however, frequently mix expressions from several theories. Consider the following examples:

- A combination of linear arithmetic and uninterpreted functions:

$$(x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge f(f(x_1) - f(x_2)) \neq f(x_3) \quad (10.1)$$

- A combination of bit-vectors and uninterpreted functions:

$$f(a[32], b[1]) = f(b[32], a[1]) \wedge a[32] = b[32] \quad (10.2)$$

- A combination of arrays and linear arithmetic:

$$x = v\{i \leftarrow e\}[j] \wedge y = v[j] \wedge x > e \wedge x > y \quad (10.3)$$

▀

In this chapter, we cover the popular **Nelson–Oppen** combination method. This method assumes that we have a decision procedure for each of the theories involved. The Nelson–Oppen combination method permits the decision procedures to communicate information with one another in a way that guarantees a sound and complete decision procedure for the combined theory.

### 10.2 Preliminaries

Let us recall several basic definitions and conventions that should be covered in any basic course on mathematical logic (see also Sect. 1.4). We assume a basic familiarity with first-order logic here.

First-order logic is a baseline for defining various restrictions thereof, which are called **theories**. It includes

- variables;
- **logical symbols** that are shared by all theories, such as the Boolean operators ( $\wedge, \vee, \dots$ ), quantifiers ( $\forall, \exists$ ) and parentheses;
- **nonlogical symbols**, namely function and predicate symbols, that are uniquely specified for each theory; and
- syntax.

It is common to consider the equality sign as a logical symbol rather than a predicate that is specific to a theory, since first-order theories without this symbol are rarely considered. We follow this convention in this chapter.

A first-order theory is defined by a set of sentences (first-order formulas in which all variables are quantified). It is common to represent such sets by a set of axioms, with the implicit meaning that the theory is the set of sentences that are derivable from these axioms. In such a case, we can talk about the “axioms of the theory”. Axioms that define a theory are called the **nonlogical axioms**, and they come in addition to the axioms that define the logical symbols, which, correspondingly, are called the **logical axioms**.

$\Sigma$

A theory is defined over a signature  $\Sigma$ , which is a set of nonlogical symbols (i.e., function and predicate symbols). If  $T$  is such a theory, we say it is a  $\Sigma$ -theory. Let  $T$  be a  $\Sigma$ -theory. A  $\Sigma$ -formula  $\varphi$  is  **$T$ -satisfiable** if there exists an interpretation that satisfies both  $\varphi$  and  $T$ . A  $\Sigma$ -formula  $\varphi$  is  **$T$ -valid**, denoted  $T \models \varphi$ , if all interpretations that satisfy  $T$  also satisfy  $\varphi$ . In other words, such a formula is  $T$ -valid if it can be derived from the  $T$  axioms and the logical axioms.

$T \models \varphi$

**Definition 10.1 (theory combination).** *Given two theories  $T_1$  and  $T_2$  with signatures  $\Sigma_1$  and  $\Sigma_2$ , respectively, the theory combination  $T_1 \oplus T_2$  is a  $(\Sigma_1 \cup \Sigma_2)$ -theory defined by the axiom set  $T_1 \cup T_2$ .*

$\oplus$

The generalization of this definition to  $n$  theories rather than two theories is straightforward.

**Definition 10.2 (the theory combination problem).** *Let  $\varphi$  be a  $\Sigma_1 \cup \Sigma_2$  formula. The theory combination problem is to decide whether  $\varphi$  is  $T_1 \oplus T_2$ -valid. Equivalently, the problem is to decide whether the following holds:*

$$T_1 \oplus T_2 \models \varphi . \quad (10.4)$$

The theory combination problem is undecidable for arbitrary theories  $T_1$  and  $T_2$ , even if  $T_1$  and  $T_2$  themselves are decidable. Under certain restrictions on the combined theories, however, the problem becomes decidable. We discuss these restrictions later on.

An important notion required in this chapter is that of a convex theory.

**Definition 10.3 (convex theory).** *A  $\Sigma$ -theory  $T$  is convex if for every conjunctive  $\Sigma$ -formula  $\varphi$*

$$\begin{aligned} (\varphi \implies \bigvee_{i=1}^n x_i = y_i) \text{ is } T\text{-valid for some finite } n > 1 &\implies \\ (\varphi \implies x_i = y_i) \text{ is } T\text{-valid for some } i \in \{1, \dots, n\}, & \end{aligned} \quad (10.5)$$

where  $x_i, y_i$ , for  $i \in \{1, \dots, n\}$ , are some variables.

In other words, in a convex theory  $T$ , if a formula  $T$ -implies a disjunction of equalities, it also  $T$ -implies at least one of these equalities separately.

**Example 10.4.** Examples of convex and nonconvex theories include:

- Linear arithmetic over  $\mathbb{R}$  is convex. A conjunction of linear arithmetic predicates defines a set of values which can be empty, a singleton, as in

$$x \leq 3 \wedge x \geq 3 \implies x = 3, \quad (10.6)$$

or infinitely large, and hence it implies an infinite disjunction. In all three cases, it fits the definition of convexity.

- Linear arithmetic over  $\mathbb{Z}$  is not convex. For example, while

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \implies (x_3 = x_1 \vee x_3 = x_2) \quad (10.7)$$

holds, neither

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \implies x_3 = x_1 \quad (10.8)$$

nor

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \implies x_3 = x_2 \quad (10.9)$$

holds.

- The conjunctive fragment of equality logic is convex. A conjunction of equalities and disequalities defines sets of variables that are equal (equality sets) and sets of variables that are different. Hence, it implies any equality between variables in the same equality set separately. Convexity follows. ▀

Many theories used in practice are in fact nonconvex, which, as we shall soon see, makes them computationally harder to combine with other theories.

## 10.3 The Nelson–Oppen Combination Procedure

### 10.3.1 Combining Convex Theories

The Nelson–Oppen combination procedure solves the theory combination problem (see Definition 10.2) for theories that comply with several restrictions.

**Definition 10.5 (Nelson–Oppen restrictions).** *In order for the Nelson–Oppen procedure to be applicable, the theories  $T_1, \dots, T_n$  should comply with the following restrictions:*

1.  $T_1, \dots, T_n$  are quantifier-free first-order theories with equality.
2. There is a decision procedure for each of the theories  $T_1, \dots, T_n$ .
3. The signatures are disjoint, i.e., for all  $1 \leq i < j \leq n$ ,  $\Sigma_i \cap \Sigma_j = \emptyset$ .
4.  $T_1, \dots, T_n$  are theories that are interpreted over an infinite domain (e.g., linear arithmetic over  $\mathbb{R}$ , but not the theory of finite-width bit vectors).

There are extensions to the basic Nelson–Oppen procedure that overcome each of these restrictions, some of which are covered in the bibliographic notes at the end of this chapter.

Algorithm 10.3.1 is the Nelson–Oppen procedure for combinations of convex theories. It accepts a formula  $\varphi$ , which must be a conjunction of literals, as input. In general, adding disjunction to a convex theory makes it nonconvex. Extensions of convex theories with disjunctions can be supported with the extension to nonconvex theories that we present later on or, alternatively, with the methods described in Chap. 11, which are based on combining a decision procedure for the theory with a SAT solver.

The first step of Algorithm 10.3.1 relies on the idea of **purification**. Purification is a satisfiability-preserving transformation of the formula, after which each atom is from a specific theory. In this case, we say that all the atoms are **pure**. More specifically, given a formula  $\varphi$ , purification generates an equisatisfiable formula  $\varphi'$  as follows:

1. Let  $\varphi' := \varphi$ .
2. For each “alien” subexpression  $\phi$  in  $\varphi'$ ,
  - (a) replace  $\phi$  with a new auxiliary variable  $a_\phi$ , and
  - (b) constrain  $\varphi'$  with  $a_\phi = \phi$ .

**Example 10.6.** Given the formula

$$\varphi := x_1 \leq f(x_1), \quad (10.10)$$

which mixes arithmetic and uninterpreted functions, purification results in

$$\varphi' := x_1 \leq a \wedge a = f(x_1). \quad (10.11)$$

In  $\varphi'$ , all atoms are pure:  $x_1 \leq a$  is an arithmetic formula, and  $a = f(x_1)$  belongs to the theory of equalities with uninterpreted functions.  $\blacksquare$

After purification, we are left with a set of pure expressions  $F_1, \dots, F_n$  such that:

$F_i$

1. For all  $i$ ,  $F_i$  belongs to theory  $T_i$  and is a conjunction of  $T_i$ -literals.
2. Shared variables are allowed, i.e., it is possible that for some  $i, j$ ,  $1 \leq i < j \leq n$ ,  $\text{vars}(F_i) \cap \text{vars}(F_j) \neq \emptyset$ .
3. The formula  $\varphi$  is satisfiable in the combined theory if and only if  $\bigwedge_{i=1}^n F_i$  is satisfiable in the combined theory.

**Algorithm 10.3.1:** NELSON–OPPEN-FOR-CONVEX-THEORIES

**Input:** A convex formula  $\varphi$  that mixes convex theories, with restrictions as specified in Definition 10.5

**Output:** “Satisfiable” if  $\varphi$  is satisfiable, and “Unsatisfiable” otherwise

1. *Purification:* Purify  $\varphi$  into  $F_1, \dots, F_n$ .
2. Apply the decision procedure for  $T_i$  to  $F_i$ . If there exists  $i$  such that  $F_i$  is unsatisfiable in  $T_i$ , return “Unsatisfiable”.
3. *Equality propagation:* If there exist  $i, j$  such that  $F_i$   $T_i$ -implies an equality between variables of  $\varphi$  that is not  $T_j$ -implied by  $F_j$ , add this equality to  $F_j$  and go to step 2.
4. Return “Satisfiable”.

**Example 10.7.** Consider the formula

$$\begin{aligned} &(f(x_1, 0) \geq x_3) \wedge (f(x_2, 0) \leq x_3) \wedge \\ &(x_1 \geq x_2) \wedge (x_2 \geq x_1) \wedge \\ &(x_3 - f(x_1, 0) \geq 1), \end{aligned} \tag{10.12}$$

which mixes linear arithmetic and uninterpreted functions. Purification results in

$$\begin{aligned} &(a_1 \geq x_3) \wedge (a_2 \leq x_3) \wedge (x_1 \geq x_2) \wedge (x_2 \geq x_1) \wedge (x_3 - a_1 \geq 1) \wedge \\ &(a_0 = 0) \wedge \\ &(a_1 = f(x_1, a_0)) \wedge \\ &(a_2 = f(x_2, a_0)). \end{aligned} \tag{10.13}$$

In fact, we applied a small optimization here, assigning both instances of the constant “0” to the same auxiliary variable  $a_0$ . Similarly, both instances of the term  $f(x_1, 0)$  have been mapped to  $a_1$  (purification, as described earlier, assigns them to separate auxiliary variables).

The top part of Table 10.1 shows the formula (10.13) divided into the two pure formulas  $F_1$  and  $F_2$ . The first is a linear arithmetic formula, whereas the second is a formula in the theory of equalities with uninterpreted functions (EUF). Neither  $F_1$  nor  $F_2$  is independently contradictory, and hence we proceed to step 3. With a decision procedure for linear arithmetic over the reals, we infer  $x_1 = x_2$  from  $F_1$ , and propagate this fact to the other theory (i.e., we add this equality to  $F_2$ ). We can now deduce  $a_1 = a_2$  in  $T_2$ , and propagate this equality to  $F_1$ . From this equality, we conclude  $a_1 = x_3$  in  $T_1$ , which is a contradiction to  $x_3 - a_1 \geq 1$  in  $T_1$ .  $\blacksquare$

**Example 10.8.** Consider the following formula, which mixes linear arithmetic and uninterpreted functions:



$F_1$ (Arithmetic over $\mathbb{R}$ )	$F_2$ (EUF)
$a_1 \geq x_3$ $a_2 \leq x_3$ $x_1 \geq x_2$ $x_2 \geq x_1$ $x_3 - a_1 \geq 1$ $a_0 = 0$	$a_1 = f(x_1, a_0)$ $a_2 = f(x_2, a_0)$
$\star x_1 = x_2$ $a_1 = a_2$ $\star a_1 = x_3$ $\star \text{FALSE}$	$x_1 = x_2$ $\star a_1 = a_2$

**Table 10.1.** Progress of the Nelson–Oppen combination procedure starting from the purified formula (10.13). The equalities beneath the middle horizontal line result from step 3 of Algorithm 10.3.1. An equality is marked with a “ $\star$ ” if it was inferred within the respective theory

$$(x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge (f(f(x_1) - f(x_2)) \neq f(x_3)) . \quad (10.14)$$

Purification results in

$$\begin{aligned}
& (x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge (f(a_1) \neq f(x_3)) \wedge \\
& (a_1 = a_2 - a_3) \wedge \\
& (a_2 = f(x_1)) \wedge \\
& (a_3 = f(x_2)) .
\end{aligned} \quad (10.15)$$

The progress of the equality propagation step, until the detection of a contradiction, is shown in Table 10.2. ▀

### 10.3.2 Combining Nonconvex Theories

Next, we consider the combination of nonconvex theories (or of convex theories together with theories that are nonconvex). First, consider the following example, which illustrates that Algorithm 10.3.1 may fail if one of the theories is not convex:

$$(1 \leq x) \wedge (x \leq 2) \wedge p(x) \wedge \neg p(1) \wedge \neg p(2) , \quad (10.16)$$

where  $x \in \mathbb{Z}$ .

Equation (10.16) mixes linear arithmetic over the integers and equalities with uninterpreted predicates. Linear arithmetic over the integers, as demonstrated in Example 10.4, is not convex. Purification results in

$$\begin{aligned}
& 1 \leq x \wedge x \leq 2 \wedge p(x) \wedge \neg p(a_1) \wedge \neg p(a_2) \wedge \\
& a_1 = 1 \wedge \\
& a_2 = 2
\end{aligned} \quad (10.17)$$

$F_1$ (arithmetic over $\mathbb{R}$ )	$F_2$ (EUF)
$x_2 \geq x_1$ $x_1 - x_3 \geq x_2$ $x_3 \geq 0$ $a_1 = a_2 - a_3$	$f(a_1) \neq f(x_3)$ $a_2 = f(x_1)$ $a_3 = f(x_2)$
$\star x_3 = 0$ $\star x_1 = x_2$ $a_2 = a_3$ $\star a_1 = 0$ $\star a_1 = x_3$	$x_1 = x_2$ $\star a_2 = a_3$  $a_1 = x_3$ FALSE

**Table 10.2.** Progress of the Nelson–Oppen combination procedure starting from the purified formula (10.15)

$F_1$ (arithmetic over $\mathbb{Z}$ )	$F_2$ (EUF)
$1 \leq x$ $x \leq 2$ $a_1 = 1$ $a_2 = 2$	$p(x)$ $\neg p(a_1)$ $\neg p(a_2)$

**Table 10.3.** The two pure formulas corresponding to (10.16) are independently satisfiable and do not imply any equalities. Hence, Algorithm 10.3.1 returns “Satisfiable”

Table 10.3 shows the partitioning of the predicates in the formula (10.17) into the two pure formulas  $F_1$  and  $F_2$ . Note that both  $F_1$  and  $F_2$  are individually satisfiable, and neither implies any equalities in its respective theory. Hence, Algorithm 10.3.1 returns “Satisfiable” even though the original formula is unsatisfiable in the combined theory.

The remedy to this problem is to consider not only implied equalities, but also implied *disjunctions* of equalities. Recall that there is a finite number of variables, and hence of equalities and disjunctions of equalities, which means that computing these implications is feasible. Given such a disjunction, the problem is split into as many parts as there are disjuncts, and the procedure is called recursively. For example, in the case of the formula (10.16),  $F_1$  implies  $x = 1 \vee x = 2$ . We can therefore split the problem into two, considering separately the case in which  $x = 1$  and the case in which  $x = 2$ . Algorithm 10.3.2 merely adds one step (step 4) to Algorithm 10.3.1: the step that performs this split.

**Algorithm 10.3.2:** NELSON–OPPEN

**Input:** A formula  $\varphi$  that mixes theories, with restrictions as specified in Definition 10.5

**Output:** “Satisfiable” if  $\varphi$  is satisfiable, and “Unsatisfiable” otherwise

1. *Purification:* Purify  $\varphi$  into  $\varphi' := F_1, \dots, F_n$ .
2. Apply the decision procedure for  $T_i$  to  $F_i$ . If there exists  $i$  such that  $F_i$  is unsatisfiable, return “Unsatisfiable”.
3. *Equality propagation:* If there exist  $i, j$  such that  $F_i$   $T_i$ -implies an equality between variables of  $\varphi$  that is not  $T_j$ -implied by  $F_j$ , add this equality to  $F_j$  and go to step 2.
4. *Splitting:* If there exists  $i$  such that
  - $F_i \implies (x_1 = y_1 \vee \dots \vee x_k = y_k)$  and
  - $\forall j \in \{1, \dots, k\}. F_i \not\Rightarrow x_j = y_j$ ,
 then apply NELSON–OPPEN recursively to

$$\varphi' \wedge x_1 = y_1, \dots, \varphi' \wedge x_k = y_k .$$

If any of these subproblems is satisfiable, return “Satisfiable”. Otherwise, return “Unsatisfiable”.

5. Return “Satisfiable”.

$F_1$ (arithmetic over $\mathbb{Z}$ )	$F_2$ (EUF)
$1 \leq x$	$p(x)$
$x \leq 2$	$\neg p(a_1)$
$a_1 = 1$	$\neg p(a_2)$
$a_2 = 2$	
$\star x = 1 \vee x = 2$	

**Table 10.4.** The disjunction of equalities  $x = a_1 \vee x = a_2$  is implied by  $F_1$ . Algorithm 10.3.2 splits the problem into the subproblems described in Tables 10.5 and 10.6, both of which return “Unsatisfiable”

**Example 10.9.** Consider the formula (10.16) again. Algorithm 10.3.2 infers ( $x = 1 \vee x = 2$ ) from  $F_1$ , and splits the problem into two subproblems, as illustrated in Tables 10.4–10.6. ▀

$F_1$ (arithmetic over $\mathbb{Z}$ )	$F_2$ (EUF)
$1 \leq x$ $x \leq 2$ $a_1 = 1$ $a_2 = 2$	$p(x)$ $\neg p(a_1)$ $\neg p(a_2)$
$x = 1$ $\star x = a_1$	$x = a_1$ FALSE

**Table 10.5.** The case  $x = a_1$  after the splitting of the problem in Table 10.4

$F_1$ (arithmetic over $\mathbb{Z}$ )	$F_2$ (EUF)
$1 \leq x$ $x \leq 2$ $a_1 = 1$ $a_2 = 2$	$p(x)$ $\neg p(a_1)$ $\neg p(a_2)$
$x = 2$ $\star x = a_2$	$x = a_2$ FALSE

**Table 10.6.** The case  $x = a_2$  after the splitting of the problem in Table 10.4

### 10.3.3 Proof of Correctness of the Nelson–Oppen Procedure

We now prove the correctness of Algorithm 10.3.1 for convex theories and for conjunctions of theory literals. The generalization to Algorithm 10.3.2 is not hard. Without proof, we rely on the fact that  $\bigwedge_i F_i$  is equisatisfiable with  $\varphi$ .

**Theorem 10.10.** *Algorithm 10.3.1 returns “Unsatisfiable” if and only if its input formula  $\varphi$  is unsatisfiable in the combined theory.*

*Proof.* Without loss of generality, we can restrict the proof to the combination of two theories  $T_1$  and  $T_2$ .

( $\Rightarrow$ , Soundness) Assume that  $\varphi$  is satisfiable in the combined theory. We are going to show that this contradicts the possibility that Algorithm 10.3.2 returns “Unsatisfiable”. Let  $\alpha$  be a satisfying assignment of  $\varphi$ . Let  $A$  be the set of auxiliary variables added as a result of the purification step (step 1). As  $\bigwedge_i F_i$  and  $\varphi$  are equisatisfiable in the combined theory, we can extend  $\alpha$  to an assignment  $\alpha'$  that includes also the variables  $A$ .

**Lemma 10.11.** *Let  $\varphi$  be satisfiable. After each loop iteration,  $\bigwedge_i F_i$  is satisfiable in the combined theory.*

*Proof.* The proof is by induction on the number of loop iterations. Denote by  $F_i^j$  the formula  $F_i$  after iteration  $j$ .

*Base case.* For  $j = 0$ , we have  $F_i^j = F_i$ , and, thus, a satisfying assignment can be constructed as described above.

*Induction step.* Assume that the claim holds up to iteration  $j$ . We shall show the correctness of the claim for iteration  $j + 1$ . For any equality  $x = y$  that is added in step 3, there exists an  $i$  such that  $F_i^j \implies x = y$  in  $T_i$ . Since  $\alpha' \models F_i^j$  in  $T_i$  by the hypothesis, clearly,  $\alpha' \models x = y$  in  $T_i$ . Since for all  $i$  it holds that  $\alpha' \models F_i^j$  in  $T_i$ , then for all  $i$  it holds that  $\alpha' \models F_i \wedge x = y$  in  $T_i$ . Hence, in step 2, the algorithm will *not* return “Unsatisfiable”.  $\blacksquare$

( $\Leftarrow$ , Completeness) First, observe that Algorithm 10.3.1 always terminates, as there are only finitely many equalities over the variables in the formula. It is left to show that the algorithm gives the answer “Unsatisfiable”. We now record a few observations about Algorithm 10.3.1. The following observation is simple to see.

$F'_i$  **Lemma 10.12.** *Let  $F'_i$  denote the formula  $F_i$  upon termination of Algorithm 10.3.1. Upon termination with the answer “Satisfiable”, any equality between  $\varphi$ ’s variables that is implied by any of the  $F'_i$  is also implied by all  $F'_j$  for any  $j$ .*

We need to show that if  $\varphi$  is unsatisfiable, Algorithm 10.3.1 returns “Unsatisfiable”. Assume falsely that it returns “Satisfiable”.

Let  $E_1, \dots, E_m$  be a set of equivalence classes of the variables in  $\varphi$  such that  $x$  and  $y$  are in the same class if and only if  $F'_1$  implies  $x = y$  in  $T_1$ . Owing to Lemma 10.12,  $x, y \in E_i$  for some  $i$  if and only if  $x = y$  is  $T_2$ -implied by  $F'_2$ .

$\Delta$  For  $i \in \{1, \dots, m\}$ , let  $r_i$  be an element of  $E_i$  (a *representative* of that set). We now define a constraint  $\Delta$  that forces all variables that are not implied to be equal to be different:

$$\Delta \doteq \bigwedge_{i \neq j} r_i \neq r_j. \quad (10.18)$$

**Lemma 10.13.** *Given that both  $T_1$  and  $T_2$  have an infinite domain and are convex,  $\Delta$  is  $T_1$ -consistent with  $F'_1$  and  $T_2$ -consistent with  $F'_2$ .*

Informally, this lemma can be shown to be correct as follows. Let  $x$  and  $y$  be two variables that are not implied to be equal. Owing to convexity, they do not have to be equal to satisfy  $F'_i$ . As the domain is infinite, there are always values left in the domain that we can choose in order to make  $x$  and  $y$  different.

Using Lemma 10.13, we argue that there are satisfying assignments  $\alpha_1$  and  $\alpha_2$  for  $F'_1 \wedge \Delta$  and  $F'_2 \wedge \Delta$  in  $T_1$  and  $T_2$ , respectively. These assignments are **maximally diverse**, i.e., any two variables that are assigned equal values by either  $\alpha_1$  or  $\alpha_2$  *must* be equal.

Given this property, it is easy to build a mapping  $M$  (an isomorphism) from domain elements to domain elements such that  $\alpha_2(x)$  is mapped to  $\alpha_1(x)$  for any variable  $x$  (this is not necessarily possible unless the assignments are maximally diverse).

As an example, let  $F_1$  be  $x = y$  and  $F_2$  be  $F(x) = G(y)$ . The only equality implied is  $x = y$ , by  $F_1$ . This equality is propagated to  $T_2$  and, thus, both  $F'_1$  and  $F'_2$  imply this equality. Possible variable assignments for  $F'_1 \wedge \Delta$  and  $F'_2 \wedge \Delta$  are

$$\begin{aligned}\alpha_1 &= \{x \mapsto \mathcal{D}_1, y \mapsto \mathcal{D}_1\} , \\ \alpha_2 &= \{x \mapsto \mathcal{D}_2, y \mapsto \mathcal{D}_2\} ,\end{aligned}\tag{10.19}$$

where  $\mathcal{D}_1$  and  $\mathcal{D}_2$  are some elements from the domain. This results in an isomorphism  $M$  such that  $M(\mathcal{D}_1) = \mathcal{D}_2$ .

Using the mapping  $M$ , we can obtain a model  $\alpha'$  for  $F'_1 \wedge F'_2$  in the combined theory by adjusting the interpretation of the symbols in  $F'_2$  appropriately. This is always possible, as  $T_1$  and  $T_2$  do not share any nonlogical symbols.

Continuing our example, we construct the following interpretation for the nonlogical symbols  $F$  and  $G$ :

$$F(\mathcal{D}_1) = \mathcal{D}_3 , \quad G(\mathcal{D}_1) = \mathcal{D}_3 .\tag{10.20}$$

As  $F'_i$  implies  $F_i$  in  $T_i$ ,  $\alpha'$  is also a model for  $F_1 \wedge F_2$  in the combined theory, which contradicts our assumption that  $\varphi$  is unsatisfiable.  $\blacksquare$

Note that without the restriction to infinite domains, Algorithm 10.3.1 may fail. The original description of the algorithm lacked such a restriction. The algorithm was later amended by adding the requirement that the theories are *stably infinite*, which is a generalization of the requirement in our presentation. The following example, given by Tinelli and Zarba in [194], demonstrates why this restriction is important.

**Example 10.14.** Let  $T_1$  be a theory over signature  $\Sigma_1 = \{f\}$ , where  $f$  is a function symbol, and axioms that enforce solutions with no more than two distinct values. Let  $T_2$  be a theory over signature  $\Sigma_2 = \{g\}$ , where  $g$  is a function symbol.

Recall that the combined theory  $T_1 \oplus T_2$  contains the union of the axioms. Hence, the solution to any formula  $\varphi \in T_1 \oplus T_2$  cannot have more than two distinct values.

Now, consider the following formula:

$$f(x_1) \neq f(x_2) \wedge g(x_1) \neq g(x_3) \wedge g(x_2) \neq g(x_3) .\tag{10.21}$$

This formula is unsatisfiable in  $T_1 \oplus T_2$  because any assignment satisfying it must use three different values for  $x_1, x_2$ , and  $x_3$ .

However, this fact is not revealed by Algorithm 10.3.2, as illustrated in Table 10.7.  $\blacksquare$

$F_1$ (a $\Sigma_1$ -formula)	$F_2$ (a $\Sigma_2$ -formula)
$f(x_1) \neq f(x_2)$	$g(x_1) \neq g(x_3)$ $g(x_2) \neq g(x_3)$

**Table 10.7.** No equalities are propagated by Algorithm 10.3.2 when checking the formula (10.21). This results in an error: although  $F_1 \wedge F_2$  is unsatisfiable, both  $F_1$  and  $F_2$  are satisfiable in their respective theories

An extension to the Nelson–Oppen combination procedure for nonstably infinite theories was given in [194], although the details of the procedure are beyond the scope of this book. The main idea is to compute, for each nonstably infinite theory  $T_i$ , a lower bound  $N_i$  on the size of the domain in which satisfiable formulas in this theory must be satisfied (it is not always possible to compute this bound). Then, the algorithm propagates this information between the theories along with the equalities. When it checks for consistency of an individual theory, it does so under the restrictions on the domain defined by the other theories.  $F_j$  is declared unsatisfiable if it does not have a solution within the bound  $N_i$  for all  $i$ .

## 10.4 Problems

**Problem 10.1 (using the Nelson–Oppen procedure).** Prove that the following formula is unsatisfiable using the Nelson–Oppen procedure, where the variables are interpreted over the integers:

$$g(f(x_1 - 2)) = x_1 + 2 \wedge g(f(x_2)) = x_2 - 2 \wedge (x_2 + 1 = x_1 - 1).$$

**Problem 10.2 (an improvement to the Nelson–Oppen procedure).**

A simple improvement to Algorithm 10.3.1 is to restrict the propagation of equalities in step 3 as follows. We call a variable *local* if it appears only in a single theory. Then, if an equality  $v_i = v_j$  is implied by  $F_i$  and not by  $F_j$ , we propagate it to  $F_j$  only if  $v_i, v_j$  are not local to  $F_i$ . Prove the correctness of this improvement.

**Problem 10.3 (proof of correctness of Algorithm 10.3.2 for the Nelson–Oppen procedure).** Prove the correctness of Algorithm 10.3.2 by generalizing the proof of Algorithm 10.3.1 given in Sect. 10.3.3.

## 10.5 Bibliographic Notes

The theory combination problem (Definition 10.2) was shown to be undecidable in [27]. The depth of the topic of combining theories resulted in an

**Aside: An Abstract Version of the Nelson–Oppen Procedure**

Let  $V$  be the set of variables used in  $F_1, \dots, F_n$ . A partition  $P$  of  $V$  induces equivalence classes, in which variables are in the same class if and only if they are in the same partition as defined by  $P$ . (Every assignment to  $V$ 's variables induces such a partition.) Denote by  $R$  the equivalence relation corresponding to these classes. The **arrangement** corresponding to  $P$  is defined by

$$ar(P) \doteq \bigwedge_{v_i R v_j, i < j} v_i = v_j \wedge \bigwedge_{\neg(v_i R v_j), i < j} v_i \neq v_j. \quad (10.22)$$

In words, the arrangement  $ar(P)$  is a conjunction of all equalities and disequalities corresponding to  $P$ , modulo reflexivity and symmetry. For example, if  $V := \{x_1, x_2, x_3\}$  and  $P := \{\{x_1, x_2\}, \{x_3\}\}$ , then

$$ar(P) := x_1 = x_2 \wedge x_1 \neq x_3 \wedge x_2 \neq x_3. \quad (10.23)$$

Now, consider the following abstract version of the Nelson–Oppen procedure:

1. Choose nondeterministically a partition  $P$  of  $V$ 's variables.
2. If one of  $F_i \wedge ar(P)$  with  $i \in \{1, \dots, n\}$  is unsatisfiable, return “Unsatisfiable”. Otherwise, return “Satisfiable”.

We have:

- *Termination.* The procedure terminates, since there is a finite number of partitions.
- *Soundness and completeness.* If the procedure returns “Unsatisfiable”, then the input formula is unsatisfiable. Indeed, if there is a satisfying assignment to the combined theory, this assignment corresponds to some arrangement; testing this arrangement leads to a termination with the result “Satisfiable”. The other direction is harder to prove, but also possible. See [193] for more details.

The nondeterministic step can be replaced with a deterministic one, by trying all such partitions possible. Hence, now it is clear that the requirement in the Nelson–Oppen procedure for sharing implied equalities can be understood as an optimization over an exhaustive search, rather than a necessity for correctness.

More generally, **abstract decision procedures** such as the one presented here are quite common in the literature. They are convenient for theoretical reasons, and can even help in designing concrete procedures in a more modular way. Abstracting some implementation details – typically by using nondeterminism – can be helpful for various reasons, such as clarity and generality, simplicity of proving an upper bound on the complexity, and simplicity of the correctness argument, as demonstrated above.



unusual history of false claims, wrong algorithms, and, correspondingly, wrong implementations in widely used tools.

The presentation of the algorithm in this chapter is based mainly on the original paper by Nelson and Oppen [137]. However, the presentation in [137] did not require that the theories were stably infinite. One year later, Oppen realized this problem and added this restriction, without fixing the proof itself [145]. A full, model-theoretic proof was provided only in 1996 by Tinelli and Harandi in [193], which also serves as a basis for the (simplified) proof in Sect. 10.3.3.

Several publications since then have extended the basic algorithms in order to combine theories with fewer restrictions. In Sect. 10.3.3, we mentioned Tinelli and Zarba's extension to the combination of nonstably infinite theories [194]. Nelson and Oppen's combination procedure in its original form, as described in this chapter, can be very inefficient. Several optimizations have been suggested, including a method for avoiding the purification step [14]. There is empirical evidence showing that the computation of the implied equalities can become a bottleneck when one is combining, for example, linear arithmetic on the basis of the Simplex method [63].

Oppen's nondeterministic combination method (see p. 237) implies a simple way to avoid equality propagation altogether. We delay a description of this idea to the next chapter (see Sect. 11.5), because its implementation is coupled with the techniques described in that chapter.

Shostak's combination procedure [179] was considered to be the major alternative to the Nelson–Oppen procedure for many years. The main difference was that it maintained a single global congruence closure data structure for all theories. The various decision procedures learned about equalities from this data structure and updated it once they had discovered new equalities. A major advantage of this method was that adding uninterpreted functions was straightforward (see Chap. 4). However, Rueß and Shankar [168] realized in 2001 that Shostak's method was in fact flawed (it was incomplete and not necessarily terminating). Any attempt to fix it turned out to be a special case of the Nelson–Oppen procedure – see, for example, the description of this matter by Barret, Dill, and Stump [14].

Krstić and Conchon showed in [108] that Shostak's method was only a way to extend decision procedures for certain theories (now called Shostak's theories) with uninterpreted-function symbols, and could not be used to combine theories. Consequently, it is not comparable with the Nelson–Oppen procedure.

In practice, the main application of the Nelson–Oppen procedure is the combination of equality logic with uninterpreted functions with other theories, for example linear arithmetic. It is implemented in this way in most state-of-the-art solvers. Note that the Nelson–Oppen procedure is not meant as a reduction technique, that is, its purpose is not to decide, for example, bit-vector arithmetic using the Simplex method.

## 10.6 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$\Sigma$	The signature of a theory, i.e., its set of nonlogical predicates and function symbols and their respective arities (i.e., those symbols that are <i>not</i> common to all first-order theories)	226
$T \models \varphi$	$\varphi$ is $T$ -valid	226
$T_1 \oplus T_2$	Denotes the theory obtained from combining the theories $T_1$ and $T_2$ , i.e., a theory over $\Sigma_1 \cup \Sigma_2$ defined by the set of axioms $T_1 \cup T_2$	226
$F_i$	The pure (theory-specific) formulas in Algorithm 10.3.1	228
$F'_i$	The formula $F_i$ upon termination of Algorithm 10.3.1	234
$\Delta$	A constraint that forces all variables that are not implied to be equal to be different	234

---

## Propositional Encodings

### 11.1 Overview

In this chapter we use several terms that were defined in Sect. 1.4, and assume the reader is familiar with the basic architecture of SAT solvers, as described in Chap. 2.

Let  $T$  be a quantifier-free first-order theory over a signature  $\Sigma$ , such that there exists a decision procedure, denoted  $DP_T$ , for the conjunctive fragment of  $T$  (i.e.,  $DP_T$  can decide a conjunction of  $\Sigma$ -literals). For example,  $T$  could be equality logic with uninterpreted functions; a possible choice for  $DP_T$  in this case is the congruence closure algorithm (Algorithm 4.1.1), described in Chap. 4. As another example,  $T$  could be disjunctive linear arithmetic and  $DP_T$  the simplex algorithm, described in Chap. 5.

In this chapter, we study a general method – a framework, really – that combines  $DP_T$  with a propositional SAT solver in various ways in order to construct a decision procedure for  $T$ . This approach has strong practical advantages, as it is both very modular and very efficient. In fact, variants of this method are considered nowadays to be the best available in terms of efficiency, modularity, and generality.<sup>1</sup> The two main engines in this framework work in tight collaboration: the SAT solver chooses those literals that need to be satisfied in order to satisfy the Boolean structure of the formula, and  $DP_T$  checks whether this choice is consistent in  $T$ .

Given a  $\Sigma$ -literal  $l$ , we associate with it a unique Boolean variable  $e(l)$ , which we call the Boolean **encoder** of this literal. Extending this idea to formulas, given a  $\Sigma$ -formula  $t$ ,  $e(t)$  denotes the Boolean formula resulting from substituting each  $\Sigma$ -literal in  $t$  with its Boolean encoder.

For example, if  $x = y$  is a  $\Sigma$ -literal, then  $e(x = y)$ , a Boolean variable, is its encoder. And if

---

<sup>1</sup> Currently (2008), this approach is considered almost as standard. All the tools that participated in the SMT-COMP tool competitions in the years 2005–2007 (see Appendix A) belong to this category.

 $e(l)$  $e(t)$

$$t := x = y \vee x = z \quad (11.1)$$

is a  $\Sigma$ -formula, then

$$e(t) := e(x = y) \vee e(x = z) . \quad (11.2)$$

For a  $\Sigma$ -formula  $t$ , the resulting Boolean formula  $e(t)$  is called the **propositional skeleton** of  $t$ .

Using this notation, we can now begin to give an overview of the methods studied in this chapter, while following a simple example. Some of the notation that we shall use in this example will be defined more formally later on.

Let  $T$  be equality logic. Given an NNF formula

$$\varphi := x = y \wedge ((y = z \wedge x \neq z) \vee x = z) , \quad (11.3)$$

we begin by computing its propositional skeleton:

$$e(\varphi) := e(x = y) \wedge (e(y = z) \wedge e(x \neq z)) \vee e(x = z) . \quad (11.4)$$

Note that since we are encoding *literals* and not *atoms*,  $e(\varphi)$  has no negations and hence is trivially satisfiable.<sup>2</sup> Let  $\mathcal{B}$  be a Boolean formula, initially set to  $e(\varphi)$ , i.e.,

$$\mathcal{B} := e(\varphi) . \quad (11.5)$$

As a second step, we pass  $\mathcal{B}$  to a SAT solver. Assume that the SAT solver returns the satisfying assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}, e(x \neq z) \mapsto \text{TRUE}, e(x = z) \mapsto \text{FALSE}\} .$$

The decision procedure  $DP_T$  now has to decide whether the conjunction of the literals corresponding to this assignment is satisfiable. We denote this conjunction by  $\hat{T}h(\alpha)$  (the “ $Th$ ” is intended to remind the reader that the result of this function is a *Theory*, and the “hat” that it is a conjunction of symbols). For the assignment above,

$$\hat{T}h(\alpha) := x = y \wedge y = z \wedge x \neq z \wedge \neg(x = z) . \quad (11.6)$$

This formula is not satisfiable, which means that the negation of this formula is a tautology. Thus  $\mathcal{B}$  is conjoined with  $e(\neg\hat{T}h(\alpha))$ , the Boolean encoding of this tautology:

$$e(\neg\hat{T}h(\alpha)) := (\neg e(x = y) \vee \neg e(y = z) \vee \neg e(x \neq z) \vee e(x = z)) . \quad (11.7)$$

This clause contradicts the current assignment, and hence *blocks* it from being repeated. Such clauses are called **blocking clauses**. In general, we denote by

<sup>2</sup> Although encoding an atom and its negation with a single variable reduces the number of encoding variables, we require here that they are encoded with two different variables. This simplifies the presentation and the proofs later on.

$t$  the formula – also called the **lemma** – returned by  $DP_T$  (in this example  $t := \neg \hat{T}h(\alpha)$ ). The negation of the current assignment is not the most effective lemma in terms of speeding up the search, as we shall see later on.


After the blocking clause has been added, the SAT solver is invoked again and suggests another assignment, for example

$$\alpha' := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}, e(x = z) \mapsto \text{TRUE}, e(x \neq z) \mapsto \text{FALSE}\}.$$

The corresponding  $\Sigma$ -formula

$$\hat{T}h(\alpha') := x = y \wedge y = z \wedge x = z \wedge \neg(x \neq z) \quad (11.8)$$

is satisfiable, which proves that  $\varphi$ , the original formula, is satisfiable. Indeed, any assignment that satisfies  $\hat{T}h(\alpha')$  also satisfies  $\varphi$ .



**Fig. 11.1.** The information exchanged between the SAT solver and a decision procedure  $DP_T$  for a conjunction of  $\Sigma$ -literals

Figure 11.1 illustrates the information flow between the two components of the decision procedure.

There are many improvements to this basic procedure, some of which we shall cover later in this chapter, and some of which are left as exercises in Sect. 11.4. One such improvement, for example, is to invoke the decision procedure  $DP_T$  after some or all *partial assignments*, rather than waiting for a full assignment. A contradicting partial assignment leads to a more powerful lemma  $t$ , as it blocks all assignments that extend it. Further, when the partial assignment is not contradictory, it can be used to derive implications that are propagated back to the SAT solver. Continuing the example above, consider the partial assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{TRUE}\}, \quad (11.9)$$

and the corresponding formula that is transferred to  $DP_T$ ,

$$\hat{T}h(\alpha) := x = y \wedge y = z. \quad (11.10)$$

This leads  $DP_T$  to conclude that  $x = z$  is implied, and hence accordingly to inform the SAT solver that  $e(x = z) \mapsto \text{TRUE}$  and  $e(x \neq z) \mapsto \text{FALSE}$  are implied by the current partial assignment  $\alpha$ . Thus, in addition to the normal Boolean constraint propagation (BCP), there is now also **theory propagation**. Such

propagation may lead to further BCP, which means that this process may iterate several times before the next decision is made by the SAT solver.

In the next few sections, we describe variations of the process demonstrated above.

## 11.2 Lazy Encodings

### 11.2.1 Definitions and Notations

As in the other chapters of this book, we focus on the satisfiability problem of NNF formulas.

Let  $lit(\varphi)$  denote the set of literals in a given NNF formula  $\varphi$ . Assuming some predefined order on the literals, we denote the  $i$ -th distinct literal in  $\varphi$  by  $lit_i(\varphi)$ .

For a given encoding  $e(\varphi)$ , we denote by  $\alpha$  an assignment, either full or partial, to the encoders in  $e(\varphi)$ . Then for an encoder  $e(lit_i)$  that is assigned a truth value by  $\alpha$ , we define the corresponding literal, denoted  $Th(lit_i, \alpha)$ , as follows:

$$Th(lit_i, \alpha) \doteq \begin{cases} lit_i & \alpha(lit_i) = \text{TRUE} \\ \neg lit_i & \alpha(lit_i) = \text{FALSE} . \end{cases} \quad (11.11)$$

Somewhat overloading the notation, we write  $Th(\alpha)$  to denote the set of literals such that their encoders are assigned by  $\alpha$ :

$$Th(\alpha) \doteq \{Th(lit_i, \alpha) \mid e(lit_i) \text{ is assigned by } \alpha\} . \quad (11.12)$$

We denote by  $\hat{Th}(\alpha)$  the conjunction over the elements in  $Th(\alpha)$ .

**Example 11.1.** To demonstrate the use of the above notation, let

$$lit_1 = (x = y), \quad lit_2 = (y = z), \quad lit_3 = (z = w) , \quad (11.13)$$

and let  $\alpha$  be a partial assignment such that

$$\alpha := \{e(lit_1) \mapsto \text{FALSE}, e(lit_2) \mapsto \text{TRUE}\} . \quad (11.14)$$

Then

$$Th(lit_1, \alpha) := \neg(x = y), \quad Th(lit_2, \alpha) := (y = z) , \quad (11.15)$$

and

$$Th(\alpha) := \{\neg(x = y), (y = z)\} . \quad (11.16)$$

Conjoining these terms gives us

$$\hat{Th}(\alpha) := \neg(x = y) \wedge (y = z) . \quad (11.17)$$

▀

### 11.2.2 A Lazy Procedure for Building Propositional Encodings

Recall that  $DP_T$  is a decision procedure for a conjunction of  $\Sigma$ -literals, where  $T$  is a theory defined over the symbols in  $\Sigma$ . Let DEDUCTION be a procedure based on  $DP_T$ , which receives a conjunction of  $\Sigma$ -literals as input, decides whether it is satisfiable, and, if the answer is negative, returns constraints over these literals, as explained below. On the basis of such a procedure, we now examine variations of the method that is demonstrated in the introduction to this chapter.

Algorithm 11.2.1 (LAZY-BASIC) decides whether a given  $\Sigma$ -formula  $\varphi$  is satisfiable. It does so by iteratively solving a propositional formula  $\mathcal{B}$ , starting from  $\mathcal{B} = e(\varphi)$ , and gradually strengthening it with encodings of constraints that are computed by DEDUCTION. (B)

In each iteration, SAT-SOLVER returns a tuple  $\langle \text{assignment}, \text{result} \rangle$  in line 4. If  $\mathcal{B}$  is unsatisfiable, then so is  $\varphi$ : LAZY-BASIC returns “Unsatisfiable” in line 5. Otherwise, LAZY-BASIC checks in line 7 whether  $\hat{T}h(\alpha)$  is satisfiable, by passing it to DEDUCTION. DEDUCTION returns a tuple of the form  $\langle \text{constraint}, \text{result} \rangle$  where the constraint is a clause over  $\Sigma$ -literals, and the result is one of {“Satisfiable”, “Unsatisfiable”}. If it is “Satisfiable”, LAZY-BASIC returns “Satisfiable” in line 8 (recall that  $\alpha$  is a full assignment). Otherwise, the clause  $t$  returned by DEDUCTION corresponds to a lemma about  $\varphi$ . In line 9, LAZY-BASIC continues by conjoining the propositional clause  $e(t)$  with  $\mathcal{B}$  and reiterating. (t)

#### Algorithm 11.2.1: LAZY-BASIC

**Input:** A formula  $\varphi$

**Output:** “Satisfiable” if  $\varphi$  is satisfiable, and “Unsatisfiable” otherwise

```

1. function LAZY-BASIC( $\varphi$ )
2.    $\mathcal{B} := e(\varphi)$ ;
3.   while (TRUE) do
4.      $\langle \alpha, res \rangle := \text{SAT-SOLVER}(\mathcal{B})$ ;
5.     if  $res = \text{“Unsatisfiable”}$  then return “Unsatisfiable”;
6.     else
7.        $\langle t, res \rangle := \text{DEDUCTION}(\hat{T}h(\alpha))$ ;
8.       if  $res = \text{“Satisfiable”}$  then return “Satisfiable”;
9.        $\mathcal{B} := \mathcal{B} \wedge e(t)$ ;

```

Consider the following three requirements from the clause  $t$  that is returned by DEDUCTION:

1. The formula  $t$  is  $T$ -valid, i.e.,  $t$  is a tautology in  $T$ . For example, if  $T$  is the theory of equality, then  $x = y \wedge y = z \implies x = z$  is  $T$ -valid.

2. The atoms in  $t$  are restricted to those appearing in  $\varphi$ .
3. The encoding of  $t$  contradicts  $\alpha$ , i.e.,  $e(t)$  is a blocking clause.

The first requirement is sufficient for guaranteeing soundness. The second and third requirements are sufficient for guaranteeing termination. Specifically, the third requirement guarantees that  $\alpha$  is not repeated.

Two of the three requirements above can be weakened, however:

- Requirement 1: the clause  $t$  can be any formula that is implied by  $\varphi$ , and not just a  $T$ -valid formula. However finding formulas that are on the one hand implied by  $\varphi$  and on the other hand fulfill the other two requirements may be as hard as deciding  $\varphi$  itself, which misses the point. In practice, the amount of effort dedicated to computing  $t$  needs to be tuned separately for each theory and decision procedure, in order to maximize the overall performance.
- Requirement 2: the clause  $t$  may refer to atoms that do not appear in  $\varphi$ , as long as the number of such new atoms is finite. For example, in equality logic, we may allow  $t$  to refer to all atoms of the form  $x_i = x_j$  where  $x_i, x_j$  are variables in  $\text{var}(\varphi)$ , even if only some of these equality predicates appear in  $\varphi$ .

### 11.2.3 Integration into DPLL

$\mathcal{B}^i$

Let  $\mathcal{B}^i$  be the formula  $\mathcal{B}$  in the  $i$ -th iteration of the loop in Algorithm 11.2.1. The constraint  $\mathcal{B}^{i+1}$  is strictly stronger than  $\mathcal{B}^i$  for all  $i \geq 1$ , because clauses are added but not removed between iterations. It is not hard to see that this implies that any conflict clause that is learned while solving  $\mathcal{B}^i$  can be reused when solving  $\mathcal{B}^j$  for  $i < j$ . This, in fact, is a special case of **incremental satisfiability**, which is supported by most modern SAT solvers.<sup>3</sup> Hence, invoking an incremental SAT solver in line 4 can increase the efficiency of the algorithm.

A better option is to integrate DEDUCTION into the DPLL-SAT algorithm, as shown in Algorithm 11.2.2. This algorithm uses a procedure ADD-CLAUSES, which adds new clauses to the current set of clauses at run time. We leave the question of why this is a better option than using an incremental SAT solver to the reader (see Problem 11.1).

### 11.2.4 Theory Propagation and the DPLL( $T$ ) Framework

Algorithm 11.2.2 can be optimized further. Consider, for example, a formula  $\varphi$  that contains an integer variable  $x_1$  and, among others, the literals  $x_1 \geq 10$  and  $x_1 < 0$ .

<sup>3</sup> Incremental satisfiability is concerned with the more general case in which clauses can also be *removed*. The question in that case is which conflict clauses can be reused. See also Problem 2.12.



**Algorithm 11.2.2:** LAZY-DPLL**Input:** A formula  $\varphi$ **Output:** “Satisfiable” if the formula is satisfiable, and “Unsatisfiable” otherwise

```

1. function LAZY-DPLL
2.   ADDCLAUSES( $cnf(e(\varphi))$ );
3.   if BCP() = “conflict” then return “Unsatisfiable”;
4.   while (TRUE) do
5.     if  $\neg$ DECIDE() then
6.        $\langle t, res \rangle :=$  DEDUCTION( $\hat{T}h(\alpha)$ );
7.       if  $res =$  “Satisfiable” then return “Satisfiable”;
8.       ADDCLAUSES( $e(t)$ );
9.       while (BCP() = “conflict”) do
10.         $backtrack-level :=$  ANALYZE-CONFLICT();
11.        if  $backtrack-level < 0$  then return “Unsatisfiable”;
12.        else BackTrack( $backtrack-level$ );
13.     else
14.       while (BCP() = “conflict”) do
15.         $backtrack-level :=$  ANALYZE-CONFLICT();
16.        if  $backtrack-level < 0$  then return “Unsatisfiable”;
17.        else BackTrack( $backtrack-level$ );

```

Assume that the DECIDE procedure assigns  $e(x_1 \geq 10) \mapsto \text{TRUE}$  and  $e(x_1 < 0) \mapsto \text{TRUE}$ . Inevitably, any call to DEDUCTION results in a contradiction between these two facts, independently of any other decisions that are made. However, Algorithm 11.2.2 does not call DEDUCTION until a full satisfying assignment is found. Thus, the time taken to complete the assignment is wasted. Moreover, as was mentioned in the introduction to this chapter, the refutation of this full assignment may be due to other reasons (i.e., a proof that a different subset of the assignment is contradictory), and, hence, additional assignments that include the same wrong assignment to  $e(x_1 \geq 10)$  and  $e(x_1 < 0)$  are not ruled out.

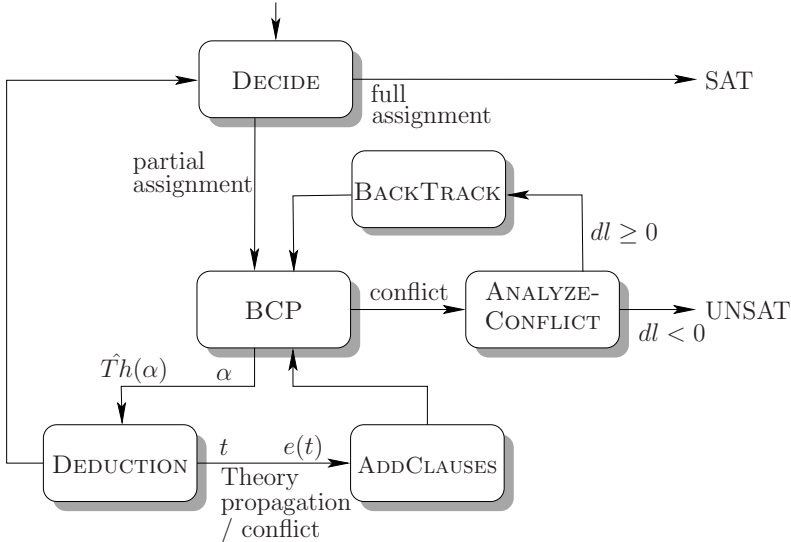
Algorithm 11.2.2 can therefore be improved by running DEDUCTION even before a full assignment to the encoders is available. This early call to DEDUCTION can serve two purposes:

1. Contradictory partial assignments are ruled out early.
2. Implications of literals that are still unassigned can be communicated back to the SAT solver, as demonstrated in the introduction to this chapter.

Continuing our example, once  $e(x_1 \geq 10)$  has been assigned TRUE, we can infer that  $e(x_1 < 0)$  must be FALSE and avoid the conflict altogether.

This brings us to the next version of the algorithm, called DPLL( $T$ ), which was first introduced in an abstract form by Tinelli [192]. As in Algo-

rithms 11.2.1 and 11.2.2, the components of the algorithm are those of DPLL and a decision procedure for a conjunctive fragment of a theory  $T$ . The name  $\text{DPLL}(T)$  emphasizes that this is a framework that can be instantiated with different theories and corresponding decision procedures.



**Fig. 11.2.** The main components of  $\text{DPLL}(T)$ . Theory propagation is implemented in **DEDUCTION**

In the version of  $\text{DPLL}(T)$  presented in Algorithm 11.2.3 (see also Fig. 11.2), **DEDUCTION** is invoked in line 11 after no more implications can be made by **BCP**. It then finds  $T$ -implied literals and communicates them to the DPLL part of the solver in the form of a constraint  $t$ .<sup>4</sup> Hence, in addition to implications in the Boolean domain, there are now also implications due to the theory  $T$ . Accordingly, this technique is known by the name **theory propagation**.

What are the restrictions on these new clauses? As before, they have to be implied by  $\varphi$  and restricted to the atoms in  $\varphi$  (or some finite superset thereof). And, as before, when  $\hat{T}h(\alpha)$  is unsatisfiable,  $e(t)$  has to block  $\alpha$ . If  $\hat{T}h(\alpha)$  is satisfiable, we require  $t$  to fulfill one of the following two conditions in order to guarantee termination:

1. The clause  $e(t)$  is an asserting clause under  $\alpha$  (asserting clauses are defined in Sect. 2.2). This implies that the addition of  $e(t)$  to  $\mathcal{B}$  and a call to **BCP** leads to an assignment to the encoder of some literal.

<sup>4</sup> **DEDUCTION** also returns the result  $res$  (whether  $\hat{T}h(\alpha)$  is satisfiable), but  $res$  is not used. We have kept it in the pseudocode in order for the algorithm to stay similar to the previous algorithms.

2. When DEDUCTION cannot find an asserting clause  $t$  as defined above,  $t$  and  $e(t)$  are equivalent to TRUE.

The second case occurs, for example, when all the Boolean variables are already assigned, and thus the formula is found to be satisfiable. In this case, the condition in line 13 is met and the procedure continues from line 5, where DECIDE is called again. Since all variables are already assigned, the procedure returns “Satisfiable”.

**Example 11.2.** Consider once again the example of the two encoders  $e(x_1 \geq 10)$  and  $e(x_1 < 0)$ . After the first of these has been set to TRUE, DEDUCTION detects that  $\neg(x_1 < 0)$  is implied, or, in other words, that

$$t := \neg(x_1 \geq 10) \vee \neg(x_1 < 0) \quad (11.18)$$

is  $T$ -valid. The corresponding encoded (asserting) clause

$$e(t) := (\neg e(x_1 \geq 10) \vee \neg e(x_1 < 0)) \quad (11.19)$$

is added to  $\mathcal{B}$ , which leads to an immediate implication of  $\neg e(x_1 < 0)$ , and possibly further implications.  $\blacksquare$

**Algorithm 11.2.3:** DPLL( $T$ )

**Input:** A formula  $\varphi$

**Output:** “Satisfiable” if the formula is satisfiable and “Unsatisfiable” otherwise

1. **function** DPLL( $T$ )
2.   ADDCLAUSES( $cnf(e(\varphi))$ );
3.   **if** BCP() = “conflict” **then return** “Unsatisfiable”;
4.   **while** (TRUE) **do**
5.     **if**  $\neg$ DECIDE() **then return** “Satisfiable”;  $\triangleright$  Full assignment
6.     **repeat**
7.       **while** (BCP() = “conflict”) **do**
8.           $backtrack\text{-}level :=$  ANALYZE-CONFLICT();
9.          **if**  $backtrack\text{-}level < 0$  **then return** “Unsatisfiable”;
10.         **else** BackTrack( $backtrack\text{-}level$ );
11.          $\langle t, res \rangle :=$  DEDUCTION( $\hat{T}h(\alpha)$ );
12.         ADDCLAUSES( $e(t)$ );
13.     **until**  $t \equiv$  TRUE

### 11.2.5 Some Implementation Details of DPLL( $T$ )

In order to improve performance, typical implementations of DPLL( $T$ ) apply various restrictions on  $t$  and how it is communicated to the DPLL-part of the procedure.<sup>5</sup>

#### Theory Propagation

Unless  $\alpha$  includes an assignment to the encoders of all theory literals, DEDUCTION performs theory propagation. Constructing an efficient mechanism for theory propagation is a challenging task in its own right. It is important to note that theory propagation is required not for *correctness*, but only for *efficiency*. Hence, the amount of effort invested in computing new implications needs to be well tuned in order to achieve the best overall performance.

The term **exhaustive theory propagation** refers to a procedure that finds and propagates *all* literals that are implied in  $T$  by  $\hat{T}h(\alpha)$ . A simple, generic way (called “plunging”) to perform theory propagation is the following. Given an unassigned theory literal  $lit_i$ , check whether  $\hat{T}h(\alpha) \wedge \neg lit_i$  is unsatisfiable. If yes, then  $lit_i$  is implied by the current assignment  $\alpha$ . The set of unassigned literals that are checked in this way depends on how exhaustive we want the theory propagation to be. This generic method is typically not the most efficient, however.

Consider, for example, the case in which  $T$  is equality logic. A simple way to perform exhaustive theory propagation in this case is the following. For each unassigned literal of the form  $x_i = x_j$ , check if the current partial assignment forms an equality path between  $x_i$  and  $x_j$ . If yes, then this literal is implied by the literals in the path. If the partial assignment forms a disequality path, the negation of this literal is implied (see Definitions 4.5 and 4.6).

In many cases exhaustive theory propagation is not “cost-effective”. In such cases, a better strategy is to perform only simple, inexpensive propagations, while ignoring more expensive ones. In the case of linear arithmetic, for example, experiments have shown that exhaustive theory propagation has a negative effect on overall performance. It is more efficient in this case to search for simple-to-find implications, such as “if  $x > c$  holds, where  $x$  is a variable and  $c$  a constant, then any literal of the form  $x > d$  is implied if  $d < c$ ”.

#### Returning Implied Assignments Instead of Clauses

Another optimization of theory propagation is concerned with the way in which the information discovered by DEDUCTION is propagated to the Boolean

---

<sup>5</sup> In addition to the optimizations and considerations described in this section, there is a detailed and more concrete description of a C++ library that implements some of these algorithms in Appendix B.

part of the solver. So far, we have required that the clause returned by DEDUCTION be  $T$ -valid. For example, if  $\alpha$  is such that  $\hat{T}h(\alpha)$  implies a literal  $lit_i$ , then

$$t := (lit_i \vee \neg \hat{T}h(\alpha)) . \quad (11.20)$$

The encoded clause  $e(t)$  is of the form

$$(e(lit_i) \vee \bigvee_{lit_j \in Th(\alpha)} \neg e(lit_j)) . \quad (11.21)$$

Nieuwenhuis, Oliveras, and Tinelli concluded that this was an inefficient method, however [142]. Their experiments on various sets of benchmarks showed that on average, fewer than 0.5% of these clauses were ever used again, and that the burden of these extra clauses slowed down the process. They suggested a better alternative, in which DEDUCTION returns a list of implied assignments (containing  $e(lit_i)$  in this case), which the SAT solver performs.

These implied assignments have no antecedent clauses in  $\mathcal{B}$ , in contrast to the standard implications due to BCP. This causes a problem in ANALYZE-CONFLICT (see Algorithm 2.2.2), which relies on antecedent clauses for deriving **conflict clauses**. As a solution, when ANALYZE-CONFLICT needs an antecedent for such an implied literal, it queries the decision procedure for an **explanation**, i.e., a clause implied by  $\varphi$  that implies this literal given the partial assignment at the time the assignment was created.

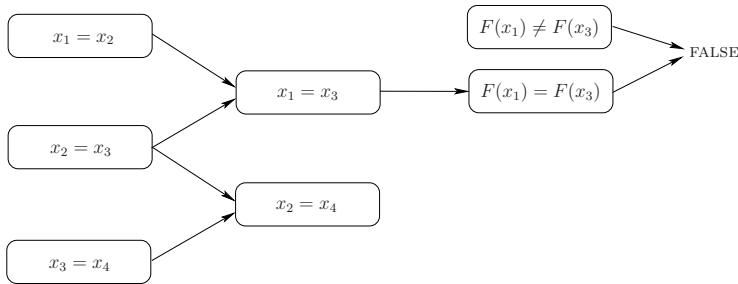
The explanation of an assignment might be the same clause that could have been delivered in the first place, but not necessarily: for efficiency reasons, typical implementations of DEDUCTION do not retain such clauses, and hence need to generate a new explanation. As an example, to explain an implied literal  $x = y$  in equality logic, one needs to search for an equality path in the equality graph between  $x$  and  $y$ , in which all the edges were present in the graph at the time that this implication was identified and propagated.

## Generating Strong Lemmas

If  $\hat{T}h(\alpha)$  is unsatisfiable, DEDUCTION returns a blocking clause  $t$  to eliminate the assignment  $\alpha$ . The stronger  $t$  is, the greater the number of inconsistent assignments it eliminates. One way of obtaining a stronger formula is to construct a clause consisting of the negation of those literals that participate in the proof of unsatisfiability of  $\hat{T}h(\alpha)$ . In other words, if  $S$  is the set of literals that serve as the premises in the proof of unsatisfiability, then the blocking clause is

$$t := \left( \bigvee_{l \in S} \neg l \right) . \quad (11.22)$$

Computing the set  $S$  corresponds to computing an *unsatisfiable core* of the formula.<sup>6</sup> Given a deductive proof of unsatisfiability, a core is easy to find. For this purpose, one may represent such a proof as a directed acyclic graph, as demonstrated in Fig. 11.3 (in this case for  $T$  being equality logic and uninterpreted functions). In this graph the nodes are labeled with literals and an edge  $(n_1, n_2)$  denotes the fact that the literal labeling node  $n_1$  was used in the inference of the literal labeling node  $n_2$ . In such a graph, there is a single sink node labeled with FALSE, and the roots are labeled with the premises (and possibly axioms) of the proof. The set of roots that can be reached by a backward traversal from the FALSE node correspond to an unsatisfiable core.



**Fig. 11.3.** The premises of a proof of unsatisfiability correspond to roots in the graph that can be reached by backward traversal from the FALSE node (in this case all roots other than  $x_3 = x_4$ ). Whereas lemmas correspond to all roots, this subset of the roots can be used for generating *strong* lemmas

## Immediate Propagation

Now consider a variation of this algorithm that calls DEDUCTION after every new assignment to an encoding variable – which may be due to either a decision or a BCP implication – rather than letting BCP finish first. Furthermore, assume that we are implementing exhaustive theory propagation as described above. This combination of features is quite common in competitive implementations of DPLL( $T$ ).

In this variant, a call to DEDUCTION *cannot lead to a conflict*, which means that it never has to return a blocking clause. A formal proof of this observation is left as an exercise (Problem 11.6). An informal justification is that if an assignment to a single encoder makes  $\hat{T}h(\alpha)$  unsatisfiable, then the negation of that assignment would have been implied and propagated in the previous

<sup>6</sup> *Unsatisfiable cores* are defined for the case of propositional CNF formulas in Sect. 2.2.6. The brief discussion here generalizes this earlier definition to inference rules other than BINARY RESOLUTION.

step by DEDUCTION. For example, if an encoder  $e(x = y)$  is implied and communicated to DEDUCTION, this literal can cause a conflict only if there is a disequality path between  $x$  and  $y$  according to the previous partial assignment. This means that in the previous step,  $\neg e(x = y)$  should have been propagated to the Boolean part of the solver.

#### Aside: Case-Splitting with BDDs

In any of the lazy algorithms described in this chapter, the service provided by the DPLL part can also be provided by a BDD. Assume we have a BDD corresponding to the propositional skeleton  $e(\varphi)$ . Each path to the “1” node in this BDD corresponds to an assignment that satisfies  $e(\varphi)$ . Hence, if one of these paths corresponds to an assignment  $\alpha$  such that  $\hat{T}h(\alpha)$  is  $T$ -satisfiable, then the original formula is satisfiable. Checking these paths one at a time is better than the basic SAT-based lazy approach for at least two reasons: first, computing each path is linear in the number of variables, in contrast to the worst-case exponential time with SAT; second, in a BDD, most of the full paths from the root to the “1” node typically do not go through all the variables, and therefore correspond to *partial* assignments, which are expected to be easier to satisfy. The drawback of this method, on the other hand, is that the BDD can become too large (recall that it may require exponential space). Some publications from the late 90’s on equality logic [87] and difference logic [130] were based on a naive version of this procedure. None of these techniques, however, apply optimizations such as strong lemmas and theory propagation, which were developed only a few years later. Such optimizations should not be too hard to implement. Theory propagation, for example, could be naturally implemented by calling DEDUCTION after visiting every node while traversing a path from top to bottom in the BDD. The formula returned by DEDUCTION should then be conjoined with the BDD, and the procedure restarted. No one, as far as we know, has experimented with a BDD-based approach combined with such optimizations.

### 11.3 Propositional Encodings with Proofs (Advanced)

In this section, we generalize the algorithms described earlier in this chapter, and in particular the process of constructing the constraint  $t$  in the procedure DEDUCTION. We assume that DEDUCTION generates deductive proofs, and show that this fact can be used to derive a tautology  $t$ , assuming that the proof system used in DEDUCTION is sound. In this method, the encoding of  $t$ , namely  $e(t)$ , represents the antecedent/consequent relations of the proof.

As a second step, we use this proof-based approach to demonstrate how to perform a full reduction from the problem of deciding  $\Sigma$ -formulas to one

of deciding propositional formulas. Such direct reductions are known by the name **eager encodings**, since, in contrast to the lazy approach, all the necessary clauses are added to the propositional skeleton up-front, or *eagerly*. The resulting propositional formula is therefore equisatisfiable with the input formula  $\varphi$ , and the SAT solver is invoked only once, with no further interaction with  $DP_T$ .

### 11.3.1 Encoding Proofs

A deductive proof is constructed using a predefined set of **proof rules** (also called *inference rules*), which we assume to be sound. A proof rule consists of a set of antecedents  $A_1, \dots, A_k$ , which are the premises that have to hold for the rule to be applicable, and a consequent  $C$ .

**Definition 11.3 (proof steps).** A proof step  $s$  is a triple  $(Rule, Conseq, Antec)$ , where *Rule* is a proof rule, *Conseq* is a proposition, and *Antec* is a (possibly empty) set of antecedents  $A_1, \dots, A_k$ .

P **Definition 11.4 (proof).** A proof  $P = \{s_1, \dots, s_n\}$  is a set of proof steps in which the transitive antecedence relation is acyclic.

The fact that the dependence between the proof steps is directed and acyclic is captured by the following definition.

**Definition 11.5 (proof graph).** A proof graph is a directed acyclic graph in which the nodes correspond to the steps, and there is an edge  $(x, y)$  if and only if the consequent of  $x$  represents an antecedent of step  $y$ .

We now define a proof step constraint.

**Definition 11.6 (proof step constraint).** Let  $s = (Rule, Conseq, Antec)$  denote a proof step, and let  $Antec = \{A_1, \dots, A_k\}$  be the set of antecedents of  $s$ . The **proof step constraint**  $psc(s)$  of  $s$  is the constraint

$$psc(s) \doteq \left( \bigwedge_{i=1}^k (A_i) \right) \implies (Conseq) . \quad (11.23)$$

We can now obtain the constraint for a whole proof by simply conjoining the constraints for all its steps.

P **Definition 11.7 (proof constraint).** Let  $P = \{s_1, \dots, s_n\}$  denote a proof. The proof constraint  $\hat{P}$  induced by  $P$  is the conjunction of the constraints induced by its steps:

$$\hat{P} \doteq \bigwedge_{i=1}^n psc(s_i) . \quad (11.24)$$



Since a proof constraint merely represents relations that are correct by construction (assuming, again, that the proof rules are sound), it is always a tautology. This, in turn, implies that DEDUCTION can safely return a proof constraint in any of the lazy algorithms described earlier in this chapter.

Blocking clauses and asserting clauses (those that are returned for the purpose of theory propagation) are special cases of proof constraints. To see why, recall that we have assumed that DEDUCTION infers these clauses through deductive proofs. But these clauses are not necessarily the proof constraints themselves. However, *there exists* a sound proof for which these clauses are the respective proof constraints. Intuitively, this is because if we infer a consequent from a set of antecedents through the application of several sound proof rules, then this means that we can construct a single sound proof rule that relates these antecedents directly to the consequent.

Using these observations, we can require DEDUCTION to return a proof constraint as defined above. Observe that if we rewrite  $p_{sc}$  (11.23) as a CNF clause, then  $e(\hat{P})$  is in CNF.

### 11.3.2 Complete Proofs

Recall that given a formula  $\varphi$ , its propositional skeleton  $e(\varphi)$  has no negations and is therefore trivially satisfiable.

**Theorem 11.8.** *If  $\varphi$  is satisfiable, then for any proof  $P$ ,  $e(\varphi) \wedge e(\hat{P})$  is satisfiable.*

Theorem 11.8 is useful if we find a proof  $P$  such that  $e(\varphi) \wedge e(\hat{P})$  is *unsatisfiable*. In such a case, the theorem implies the unsatisfiability of  $\varphi$ . In other words, we would like to restrict ourselves to proofs with the following property:

**Definition 11.9 (complete proof).** *A proof  $P$  is called complete with respect to  $\varphi$  if  $e(\varphi) \wedge e(\hat{P})$  is equisatisfiable with  $\varphi$ .*

Note that Theorem 11.8 implies that if the formula is satisfiable, then any proof is complete. Our focus is therefore on unsatisfiable formulas.

**Theorem 11.10.** *Given a sound and complete deductive decision procedure for a conjunction of  $\Sigma$ -literals, there is an algorithm for deriving a complete proof for every  $\Sigma$ -formula.*

*Proof. (sketch)* Let  $\varphi'$  be the DNF representation of a  $\Sigma$ -formula  $\varphi$ . Let  $DP_T$  be a deductive, sound, and complete decision procedure for a conjunction of  $\Sigma$ -literals. We use  $DP_T$  to prove each of the terms in  $\varphi'$ . The union of the proof steps in these proofs (together with a proof step for case-splitting) constitutes a complete proof for  $\varphi'$ . ▀

The goal, however, is to find complete proofs with smaller practical complexity than that of performing such splits: there is no point in having a procedure in which the encoding process is as complex as performing the proof directly.

Our strategy is to find deductive proofs that begin from the literals of the input formulas, leaving it for the SAT solver to deal with the Boolean structure.

**Example 11.11.** Consider the unsatisfiable formula

$$\varphi := x = 5 \wedge (x < 0 \vee x \neq 5) . \tag{11.25}$$

The skeleton of  $\varphi$  is

$$e(\varphi) := e(x = 5) \wedge (e(x < 0) \vee e(x \neq 5)) . \tag{11.26}$$

$$\begin{array}{l} \frac{}{a < succ^i(a)} \text{ (ORDERING I)} \qquad \frac{x < y \quad y < x}{\text{FALSE}} \text{ (ORDERING II)} \\ \frac{x \neq x}{\text{FALSE}} \text{ (EQ-CONTRADICTION)} \qquad \frac{x = a \quad P}{P[x/a]} \text{ (SUBSTITUTION)} \end{array}$$

**Fig. 11.4.** Inference rules for the proof  $P$  in Fig. 11.5. ORDERING I is an axiom schema, which uses  $succ^i(a)$  to denote the  $i$ -th successor,  $i > 0$ , of  $a$

<i>Consequent Rule</i>			$e(psc(s))$
1.	$x = 5$	Premise	
2.	$x \neq 5$	Premise	
3.	$x < 0$	Premise	
4.	$5 < 0$	SUBSTITUTION, 1, 3	$e(x = 5) \wedge e(x < 0) \implies e(5 < 0)$
5.	$0 < 5$	ORDERING I ( $i = 5$ )	$e(0 < 5)$
6.	FALSE	ORDERING II, 4, 5	$e(5 < 0) \wedge e(0 < 5) \implies \text{FALSE}$
7.	$5 \neq 5$	SUBSTITUTION, 1, 2	$e(x = 5) \wedge e(x \neq 5) \implies e(5 \neq 5)$
8.	FALSE	EQ-CONTRADICTION, 7	$e(5 \neq 5) \implies \text{FALSE}$

**Fig. 11.5.** Proof of unsatisfiability of  $\varphi := x = 5 \wedge (x < 0 \vee x \neq 5)$ , using the rules in Fig. 11.4. The only premises are the literals in the formula. The proof steps are annotated in the right column with the constraints that they induce

Using the proof rules in Fig. 11.4, we show a contradiction using the proof  $P$ , which appears in Fig. 11.5. Note that  $P$  uses only literals as antecedents. The encoded proof constraint is:

$$\begin{aligned}
e(\hat{P}) &:= (e(x = 5) \wedge e(x < 0) \implies e(5 < 0)) \\
&\quad \wedge (e(0 < 5)) \\
&\quad \wedge (e(5 < 0) \wedge e(0 < 5) \implies \text{FALSE}) \\
&\quad \wedge (e(x = 5) \wedge e(x \neq 5) \implies e(5 \neq 5)) \\
&\quad \wedge (e(5 \neq 5) \implies \text{FALSE}) .
\end{aligned} \tag{11.27}$$

The conjunction of  $e(\varphi)$  and  $e(\hat{P})$  is unsatisfiable, and thus, owing to Theorem 11.8,  $\varphi$  is unsatisfiable.  $\blacksquare$

How can we find such proofs that use only the literals of  $\varphi$  as premises? The next subsection answers this question. It introduces the *eager* approach, in which all necessary proof steps, starting from  $\varphi$ 's literals as premises, are computed a priori. The resulting proof constraint, conjoined with the propositional skeleton, are then sent to a SAT solver.

### 11.3.3 Eager Encodings

All the algorithms that we have seen so far can be interpreted as being aimed at constructing a propositional formula  $\mathcal{B}$  that is equisatisfiable with the original formula  $T$ . If  $\mathcal{B}$  has this property, we say that it is a **propositional encoding** of  $\varphi$ . The various lazy approaches that we studied in Sect. 11.2 can be understood as building a propositional encoding of  $\varphi$  *incrementally* (or “*lazily*”).

In contrast, Algorithm 11.3.1 (EAGER-ENCODING) computes a propositional encoding of a given formula  $\varphi$  in a single step. All the proof steps that might be necessary are assumed to be performed by the DEDUCTION procedure before the propositional engine is called.<sup>7</sup> The premises that DEDUCTION can use either are axioms or belong to the set that it receives as input.

#### Algorithm 11.3.1: EAGER-ENCODING

**Input:** A formula  $\varphi$

**Output:** “Satisfiable” if  $\varphi$  is satisfiable and “Unsatisfiable” otherwise

1. **function** EAGER-ENCODING( $\varphi$ )
2.    $P :=$  DEDUCTION( $lit(\varphi)$ );
3.    $\varphi_E := e(\varphi) \wedge e(\hat{P})$ ;
4.    $\langle \alpha, res \rangle :=$  SAT-SOLVER( $\varphi_E$ );
5.   **if**  $res =$  “Unsatisfiable” **then return** “Unsatisfiable”;
6.   **else return** “Satisfiable”;

<sup>7</sup> We overload DEDUCTION in this algorithm: it receives a set of literals rather than their conjunction as input, and returns a proof rather than an arbitrary tautology over the atoms of the input.

The SAT-SOLVER procedure returns a pair  $\langle \text{assignment}, \text{result} \rangle$  in line 4, where “result” is one of {“Satisfiable”, “Unsatisfiable”}, and “assignment” is a full satisfying assignment if “result” = “Satisfiable”, and undefined otherwise.

The result of EAGER-ENCODING matches the result returned by SAT-SOLVER, as  $\varphi$  and  $\varphi_E$  are equisatisfiable. It is left for us to describe sufficient conditions for complete proofs. In other words, it is enough to prove that a given implementation of DEDUCTION fulfills any one of these conditions in order to establish completeness of the procedure.

### 11.3.4 Criteria for Complete Proofs

Let  $\alpha$  be either a partial or a full truth assignment to  $e(\varphi)$ . The following notation is used: we write  $Th(\alpha) \longrightarrow_P \text{FALSE}$  if  $P$  leads to FALSE using  $Th(\alpha)$  as premises. The following example demonstrates the use of this notation.

**Example 11.12.** Let

$$lit_1 := x_1 > x_2, lit_2 := x_2 \leq x_1 \quad (11.28)$$

be the literals of a formula  $\varphi$ . Now consider the assignment

$$\alpha := \{e(x_1 > x_2) \mapsto \text{TRUE}, e(x_2 \leq x_1) \mapsto \text{FALSE}\}. \quad (11.29)$$

Thus, we have

$$Th_1(\alpha) := x_1 > x_2, Th_2(\alpha) := \neg(x_2 \leq x_1), \quad (11.30)$$

and

$$Th(\alpha) := \{x_1 > x_2, x_2 > x_1\}. \quad (11.31)$$

Now consider the proof rules

$$\frac{x_i > x_j \quad x_j > x_k}{x_i > x_k} \text{ (>-TRANS)}, \quad (11.32)$$

$$\frac{x_i > x_i}{\text{FALSE}} \text{ (>-CONTR)},$$

and consider  $Th(\alpha)$  as the set of premises. Let  $P$  be the following proof:

$$P := \{ (>-TRANS, x_1 > x_1, Th(\alpha)), (>-CONTR, \text{FALSE}, x_1 > x_1) \}. \quad (11.33)$$

This proof shows that  $Th(\alpha)$  is inconsistent, i.e.,  $Th(\alpha) \longrightarrow_P \text{FALSE}$ .  $\blacksquare$

The following theorem defines a sufficient condition for the completeness of a proof.

**Theorem 11.13 (sufficient condition #1 for completeness).** *Let  $\varphi$  be an unsatisfiable formula. A proof  $P$  is complete with respect to  $\varphi$  if, for every full assignment  $\alpha$  to  $e(\varphi)$ ,*

$$\alpha \models e(\varphi) \implies Th(\alpha) \longrightarrow_P \text{FALSE}. \quad (11.34)$$

The premise of this theorem can be weakened, however, which leads to a stronger theorem. We need the following definitions.

**Definition 11.14 (satisfying partial assignment).** A partial assignment  $\alpha$  to the variables in  $\text{var}(e(\varphi))$  satisfies  $e(\varphi)$  if, for any full assignment  $\alpha'$  that extends  $\alpha$ ,  $\alpha' \models e(\varphi)$ .

**Definition 11.15 (minimal satisfying assignment).** An assignment  $\alpha$  (either full or partial) that satisfies  $e(\varphi)$  is called minimal if, for any  $e \in \text{var}(e(\varphi))$  that is assigned by  $\alpha$ ,  $\alpha$  without  $e$  is not a satisfying partial assignment to  $e(\varphi)$ .

**Theorem 11.16 (sufficient condition #2 for completeness).** Let  $\varphi$  be an unsatisfiable formula, and let  $A$  denote the set of minimal satisfying assignments of  $e(\varphi)$ . A proof  $P$  is complete with respect to  $\varphi$  if, for every  $\alpha \in A$ ,  $\text{Th}(\alpha) \xrightarrow{P} \text{FALSE}$ .

Now consider a weaker requirement for complete proofs.

**Theorem 11.17 (sufficient condition #3 for completeness).** Let  $\varphi$  be an unsatisfiable formula, and let  $A$  denote the set of minimal satisfying assignments of  $e(\varphi)$ . A proof  $P$  is complete with respect to  $\varphi$  if, for every  $\alpha \in A$  and for some unsatisfiable core  $\text{Th}^{uc}(\alpha) \subseteq \text{Th}(\alpha)$ ,  $\text{Th}^{uc}(\alpha) \xrightarrow{P} \text{FALSE}$ .

Note that there is at least one unsatisfiable core because  $\hat{\text{Th}}(\alpha)$  must be unsatisfiable if  $\alpha \models e(\varphi)$  and  $\varphi$  is unsatisfiable.

It is not hard to see that Theorem 11.17 implies Theorem 11.16, which, in turn, implies Theorem 11.13 (see Problem 11.8). Hence we shall prove only Theorem 11.17.

*Proof.* Let  $\varphi$  be an unsatisfiable formula. Assume falsely that  $e(\varphi) \wedge e(\hat{P})$  is satisfiable, where  $P$  satisfies the premise of Theorem 11.17, i.e., for each minimal satisfying assignment  $\alpha$  of  $e(\varphi)$ , it holds that  $\text{Th}^{uc}(\alpha) \xrightarrow{P} \text{FALSE}$  for some unsatisfiable core  $\text{Th}^{uc}(\alpha) \subseteq \text{Th}(\alpha)$ . Let  $\alpha'$  be the satisfying assignment, and let  $\alpha$  be a minimal satisfying assignment of  $e(\varphi)$  that can be extended to  $\alpha'$ . Let  $\text{Th}^{uc}(\alpha) \subseteq \text{Th}(\alpha)$  denote an unsatisfiable core of  $\text{Th}(\alpha)$  such that  $\text{Th}^{uc}(\alpha) \xrightarrow{P} \text{FALSE}$ . This implies that  $e(\hat{P})$  evaluates to FALSE when the encoders of the literals in this core are evaluated according to  $\alpha$ . This implies that  $e(\varphi) \wedge e(\hat{P})$  evaluates to FALSE under  $\alpha$  – a contradiction.  $\blacksquare$

The problem, now, is to find a proof  $P$  that fulfills one or more of these sufficient conditions.

### 11.3.5 Algorithms for Generating Complete Proofs

Recall that by Theorem 11.10 (or, rather, by its proof), a sound and complete deductive decision procedure for a conjunction of terms can be used to

generate complete proofs, simply by case-splitting and conjoining the proof steps. As discussed earlier, however, this type of procedure misses the point, as we want to find such proofs with less effort than if we were to use splitting. We now study strategies for modifying such procedures so that they generate complete proofs from disjunctive formulas with potentially less effort than that required by splitting. The procedures that we study in this section are generic, and fulfill conditions much stronger than are required according to Theorems 11.13, 11.16, and 11.17. More specific procedures are expected to be more efficient and utilize the weaker conditions in those theorems.

We need the following definition.

$\Gamma$  **Definition 11.18 (saturation).** *Let  $\Gamma$  be an inference system (i.e., a set of inference rules and axioms, including schemas). We say that the process of applying  $\Gamma$  to a set of premises saturates if no new consequents can be derived on the basis of these premises and previously derived consequents.  $\Gamma$  is said to be saturating if the process of applying it to any finite set of premises saturates.*

In this section, we consider the class of decision procedures whose underlying inference system is saturating. Many popular decision procedures belong to this class. For example, the simplex method, the Fourier–Motzkin elimination and the Omega test, all of which are covered in Chap. 5, can be presented as being based on deduction and belong to this class.<sup>8</sup>

As before, let  $DP_T$  be a deductive decision procedure in this class for conjunction of  $\Sigma$ -literals, and let  $\Gamma$  be the set of inference rules that it can use. Let  $\varphi$  be a (disjunctive)  $\Sigma$ -formula. Now consider the following procedure:

Apply the rules in  $\Gamma$  to  $lit(\varphi)$  until saturation.

Since every inference that is possible after case-splitting is also possible here, this procedure clearly generates a complete proof. Note that the generality of this variant comes at the price of completely ignoring the inference strategy applied by the original decision procedure  $DP_T$ , which entails a sacrifice in efficiency. Nevertheless, even with this general scheme, the number of inferences is expected to be much smaller than that obtained using case-splitting, because the same inference is never repeated (whereas it can be repeated an exponential number of times with case-splitting).

Specific decision procedures that belong to this class can be changed in a way that results in a more efficient procedure, however. Here, we consider the case of projection-based decision procedures, and present it through an example, namely the Fourier–Motzkin procedure for linear arithmetic (see Sect. 5.4).

<sup>8</sup> It is not so simple to present the simplex method as a deductive system, but such a presentation appears in the literature. See Nelson [134] and Ruesch and Shankar [170] for a deductive version of the simplex method.

The Fourier–Motzkin procedure, although not presented this way in Chap. 5, can be reformulated as a deductive system, by applying the following rules:

$$\frac{UB \geq x \quad x \geq LB}{UB \geq LB} \quad (\text{PROJECT}) \quad (11.35)$$

(where  $UB$  and  $LB$  are linear constraints that do not include  $x$ ), and, for any two constants  $l, u$  such that  $l \leq u$ ,<sup>9</sup>

$$\frac{l > u}{\text{FALSE}} \quad (\text{CONSTANTS}) . \quad (11.36)$$

Given a conjunction of normalized linear arithmetic predicates  $\phi$  (i.e., equalities and negations are eliminated, as explained in Sect. 5.4), the strategy of the Fourier–Motzkin procedure can be reformulated, informally, as follows:

1. If  $\text{var}(\phi) = \emptyset$  return “Satisfiable”.
2. Choose a variable  $v \in \text{var}(\phi)$ .
3. For every upper bound  $UB$  and a lower bound  $LB$  on  $x$ , apply the rule PROJECT.
4. Simplify the resulting constraints by accumulating the coefficients of the same variable.
5. Remove all the constraints that contain  $x$ .
6. If the rule CONSTANTS is applicable, return “Unsatisfiable”.
7. Go to step 2.

Now consider the following variation of this procedure, which is meant for generating complete proofs rather than for deciding a given formula. Replace step 6 with

6. If the rule CONSTANTS is applicable, apply it.

The following example demonstrates this technique.

**Example 11.19.** Consider the following formula,

$$\varphi := (2x_1 - x_2 \leq 0) \wedge (x_3 - x_1 \leq -1) \wedge ((1 \leq x_3) \wedge (x_2 \leq 3)) \vee ((0 \leq x_3) \wedge (x_2 \leq 1)) \quad (11.37)$$

and its corresponding skeleton,

$$e(\varphi) := e(2x_1 - x_2 \leq 0) \wedge e(x_3 - x_1 \leq -1) \wedge ((e(1 \leq x_3) \wedge e(x_2 \leq 3)) \vee (e(0 \leq x_3) \wedge e(x_2 \leq 1))) . \quad (11.38)$$

Let  $x_1, x_2, x_3$  be the elimination order. The corresponding proof, according to the newly suggested procedure, is

---

<sup>9</sup> This means that the rule is applicable only when this condition is met. Such conditions are called **side conditions**.

$$\begin{aligned}
P := \{ & (\text{PROJECT} \quad , 2x_3 - x_2 \leq -2 , \{2x_1 - x_2 \leq 0, x_3 - x_1 \leq -1\} ) \\
& (\text{PROJECT} \quad , 2x_3 \leq 1 \quad , \{x_2 \leq x_3, 2x_3 - x_2 \leq -2\} ) \\
& (\text{PROJECT} \quad , 2x_3 \leq -1 \quad , \{x_2 \leq 1, 2x_3 - x_2 \leq -2\} ) \\
& (\text{PROJECT} \quad , 1 \leq \frac{1}{2} \quad , \{1 \leq x_3, 2x_3 \leq 1\} ) \\
& (\text{CONSTANTS} , \text{FALSE} \quad , \{1 \leq \frac{1}{2}\} ) \\
& (\text{PROJECT} \quad , 1 \leq -\frac{1}{2} \quad , \{1 \leq x_3, 2x_3 \leq -1\} ) \\
& (\text{CONSTANTS} , \text{FALSE} \quad , \{1 \leq -\frac{1}{2}\} ) \\
& (\text{PROJECT} \quad , 0 \leq \frac{1}{2} \quad , \{0 \leq x_3, 2x_3 \leq 1\} ) \\
& (\text{PROJECT} \quad , 0 \leq -\frac{1}{2} \quad , \{0 \leq x_3, 2x_3 \leq -1\} ) \\
& (\text{CONSTANTS} , \text{FALSE} \quad , \{0 \leq -\frac{1}{2}\} ) \} .
\end{aligned} \tag{11.39}$$

The corresponding encoding of the proof constraint is thus

$$\begin{aligned}
e(\hat{P}) := & (e(2x_1 - x_2 \leq 0) \wedge e(x_3 - x_1 \leq -1) \implies e(2x_3 - x_2 \leq -2)) \\
& \wedge (e(x_2 \leq x_3) \wedge e(2x_3 - x_2 \leq -2) \implies e(2x_3 \leq 1)) \\
& \wedge (e(x_2 \leq 1) \wedge e(2x_3 - x_2 \leq -2) \implies e(2x_3 \leq -1)) \\
& \wedge (e(1 \leq x_3) \wedge e(2x_3 \leq 1) \implies e(1 \leq \frac{1}{2})) \\
& \wedge (e(1 \leq \frac{1}{2}) \implies \text{FALSE}) \\
& \wedge (e(1 \leq x_3) \wedge e(2x_3 \leq -1) \implies e(1 \leq -\frac{1}{2})) \\
& \wedge (e(1 \leq -\frac{1}{2}) \implies \text{FALSE}) \\
& \wedge (e(0 \leq x_3) \wedge e(2x_3 \leq 1) \implies e(0 \leq \frac{1}{2})) \\
& \wedge (e(0 \leq x_3) \wedge e(2x_3 \leq -1) \implies e(0 \leq -\frac{1}{2})) \\
& \wedge (e(0 \leq -\frac{1}{2}) \implies \text{FALSE}) .
\end{aligned} \tag{11.40}$$

The conjunction of (11.38) and (11.40) is unsatisfiable, as is the original formula  $\varphi$ .

This example demonstrates also the disadvantage of this approach in comparison with the lazy approach: many of the added constraints are redundant. In this example,  $e(1 \leq x_3)$  and  $e(x_2 \leq x_3)$  do not have to be satisfied simultaneously with  $e(0 \leq x_3)$  and  $e(x_2 \leq 1)$ , because of the disjunction between them. Hence a constraint such as  $(e(1 \leq x_3) \wedge e(2x_3 \leq -1) \implies e(1 \leq -\frac{1}{2}))$  is redundant, because  $e(2x_3 \leq -1)$  is forced to be TRUE only when  $e(x_2 \leq 1)$  is assigned TRUE. Hence,  $e(1 \leq -\frac{1}{2})$  is assigned TRUE only when at least  $e(1 \leq x_3)$  and  $e(x_2 \leq 1)$  are assigned TRUE, whereas we have seen that these two encoders need not be satisfied simultaneously in order to satisfy  $e(\varphi)$ . ■

Two questions come to mind. First, does the above procedure generate fewer proof steps than does saturation? The answer is yes. To see why, consider what a saturation-based procedure would do on the basis of the above two rules. For each variable, at each step, it would apply the rule PROJECT. Hence, the overall set of proof steps corresponds to the union of proof steps when the Fourier–Motzkin procedure is applied in all possible orders. Second, is the generated proof still complete? Again, the answer is yes, and the proof is based on showing that it maintains the premise of Theorem 11.13. In fact, it maintains a much stronger condition – see Problem 11.9.



## 11.4 Problems

**Problem 11.1 (incrementality in LAZY-DPLL).** Recall that an incremental SAT solver is one that knows which conflict clauses can be reused when given a problem similar to the previous one (i.e., some clauses are added and others are erased). Is there a difference between Algorithm 11.2.2 (LAZY-DPLL) and replacing line 4 in Algorithm 11.2.1 with a call to an incremental SAT solver?

**Problem 11.2 (an optimization for Algorithms 11.2.1–11.2.3?).**

1. Consider the following variation of Algorithms 11.2.1–11.2.3 for an input formula  $\varphi$  given in NNF. Rather than sending  $\hat{T}h(\alpha)$  to DEDUCTION, send  $\bigwedge Th_i$  for all  $i$  such that  $\alpha(e_i) = \text{TRUE}$ . For example, given an assignment

$$\alpha := \{e(x = y) \mapsto \text{TRUE}, e(y = z) \mapsto \text{FALSE}, e(x = z) \mapsto \text{TRUE}\}, \quad (11.41)$$

check

$$x = y \wedge x = z. \quad (11.42)$$

Is this variation correct? Prove that it is correct or give a counterexample.

2. Show an example in which the above variation reduces the number of iterations between DEDUCTION and the SAT solver.

**Problem 11.3 (theory propagation for difference logic).** Suggest an efficient procedure that performs exhaustive theory propagation for the case of difference logic (difference logic is presented in Sect. 5.7).

**Problem 11.4 (theory propagation).** Let  $DP_T$  be a decision procedure for a conjunction of  $\Sigma$ -literals. Suggest a procedure for performing exhaustive theory propagation with  $DP_T$ .

**Problem 11.5 (pseudocode for a variant of DPLL( $T$ )).** Recall the variant of DPLL( $T$ ) suggested at the end of Sect. 11.2.5, where the partial assignment is sent to the theory solver after every assignment to an encoder, rather than only after BCP. Write pseudocode for this algorithm, and a corresponding drawing in the style of Fig. 11.2.

**Problem 11.6 (exhaustive theory propagation).** It was claimed in Sect. 11.2.5 that with exhaustive theory propagation, conflicts cannot occur in DEDUCTION and that, consequently, DEDUCTION never returns blocking clauses. Prove this claim.

**Problem 11.7 (practicing eager encodings).** Consider the following formula:

$$\varphi := (2x_1 - x_2 \leq 0) \wedge ((2x_2 - 4x_3 \leq 0) \vee (x_3 - x_1 \leq -1) \vee ((0 \leq x_3) \wedge (x_2 \leq 1))). \quad (11.43)$$

Show an eager encoding of this formula, using the rules PROJECT and CONSTANTS (see p. 261). Check that the resulting formula is equisatisfiable with  $\varphi$ .

**Problem 11.8 (proof of Theorems 11.13 and 11.16).** Prove Theorems 11.13 and 11.16 *without* referring to Theorem 11.17.

**Problem 11.9 (complete proofs).** Consider the variant of the Fourier–Motzkin procedure that was presented in Sect. 11.3.5. Show that the generated proof  $P$  proves the inconsistency of every inconsistent subset of literals. In what sense does this fulfill a stronger requirement than what is required by Theorem 11.13?

**Problem 11.10 (complexity of eager encoding with the Fourier–Motzkin procedure).** Consider the variant of the Fourier–Motzkin procedure that was presented in Sect. 11.3.5. What is the complexity of this decision procedure?

## 11.5 Bibliographic Notes

The following are some bibliographic details about the development of the lazy and the eager encoding frameworks.

### Lazy Encodings

Alessandro Armando, Claudio Castellini and Enrico Giunchiglia in [4] proposed a solver based on an interplay between a SAT solver and a theory solver, in a fashion similar to the simple lazy approach introduced at the beginning of this chapter in 1999. Their solver was tailored to a single theory called disjunctive temporal constraints, which is a restricted version of difference logic. In fact, they combined lazy with eager reasoning: they used a preprocessing step that adds a large set of constraints to the propositional skeleton (constraints of the form  $(\neg e_1 \vee \neg e_2)$  if a preliminary check discovers that the theory literals corresponding to these encoders contradict each other), which saves a lot of work later for the lazy-style engine. In the same year LPSAT [202] was introduced, which also includes many of the features described in this chapter, including a process of learning strong lemmas.

The basic idea of integrating DPLL with a decision procedure for some (single) theory was suggested even earlier than that mostly in the domain of modal and description logics [5, 86, 97, 148].

The major progress in efficient SAT solving due to the Chaff SAT solver in 2001 [133], led several groups, a year later, to (independently) propose decision procedures that leverage this progress, all of which correspond to some variation of the lazy approach described in Sect. 11.2: CVC [13, 188] by Aaron

Stump, Clark Barrett and David Dill; ICS-SAT [74] by Jean-Christophe Filliatre, Sam Owre, Herald Ruess and Natarajan Shankar; MATHSAT [6] by Gilles Audemard, Piergiorgio Bertoli, Alessandro Cimatti, Artur Kornilowicz, and Roberto Sebastiani; DLSAT [120] by Moez Mahfoudh, Peter Niebert, Eugene Asarin, and Oded Maler; and VERIFUN [76] by Cormac Flanagan, Rajeev Joshi, Xinming Ou and Jim Saxe. Most of these tools were built as generic engines that can be extended with different decision procedures. Since the introduction of these tools, this approach has become mainstream, and at least ten other solvers based on the same principles have been developed and published. In fact, all the tools that participated in the SMT-COMP competitions in 2005–2007 (see Appendix A) belong to this category of solvers.

DPLL( $T$ ) was originally described in abstract terms, in the form of a calculus, by Cesare Tinelli in [192]. Theory propagation had already appeared under various names in the papers by Armando et al. [4] and Audemard et al. [6] mentioned above. Efficient theory propagation tailored to the underlying theory  $T$  ( $T$  being EUF in that case) appeared first in a paper by Ganzinger et al. [79]. These authors also introduced the idea of propagating theory implications by maintaining a stack of such implied assignments, coupled with the ability to explain them a posteriori, rather than sending asserting clauses to the DPLL part of the solver. The idea of minimizing the lemmas (blocking clauses) can be attributed to Leonardo de Moura and Herald Ruess [60], although, as we mentioned earlier, finding small lemmas already appeared in the description of LPSAT.

Various details of how a DPLL-based SAT solver could be transformed into a DPLL( $T$ ) solver were described for the case of EUF in [79] and for difference logic in [140]. A good description of DPLL( $T$ ), starting from an abstract DPLL procedure and ending with fine details of implementation, was given in [142]. A very comprehensive survey on lazy SMT was given by Sebastiani [175]. There has been quite a lot of research on how to design  $T$ -solvers that can give *explanations*, which, as pointed out in Sect. 11.2.5, is a necessary component for efficient implementation of this framework – see, for example, [62, 141, 190].

Among the new generation of tools, let us mention four. CVC-LITE [9] and later CVC-3[11], the development of which was led by Clark Barrett, are modernized versions of CVC, which extend it with new theories (such as an extensive support for recursive data-types), improve its implementation of various decision procedures, enable each theory to produce a proof that can be checked independently with an external tool, make it compatible with the SMT-LIB standard (see Appendix A), and so forth. BARCELOGIC [79] was developed by Robert Nieuwenhuis and Albert Oliveras, and won the SMT-COMP 2005 competition. Finally, YICES, which was developed by Bruno Dutertre and Leonardo de Moura, won in most of the categories in both SMT-COMP 2006 and SMT-COMP 2007. Only a few details of YICES [58] have been published. It is a DPLL( $T$ ) solver, which uses very efficient implementations of decision procedures for the various theories it supports. Its decision procedure for lin-

ear arithmetic, based on the generalized simplex method, is the best known for 2006–2007 and has been described in [70]. Finally, let us also mention the Decision Procedure Toolkit (DPT), released as open source by Intel<sup>TM</sup>, which combines a modern implementation of DPLL(T) with various theory solvers. DPT was written by Amit Goel, Jim Grundy and Sava Krstic and is described in [109].

The lazy approach opens up new opportunities with regard to implementing the Nelson–Oppen combination procedure, described in the previous chapter. A contribution by Bozzano et al. [28] suggests a technique called **delayed theory combination**. Each pair of shared variables is encoded with a new Boolean variable (resulting in a quadratic increase in the number of variables). After all the other encoding variables have been assigned, the SAT solver begins to assign values (arbitrary at first) to the new variables, and continues as usual, i.e., after every such assignment, the current partial assignment is sent to a theory solver. If any one of the theory solvers “objects” to the arrangement implied by this assignment (i.e., it finds a conflict with the current assignment to the other literals), this leads to a conflict and backtracking. Otherwise, the formula is declared satisfiable. This way, each theory can be solved separately, without passing information about equalities. Empirically, this method is very effective, both because the individual theory solvers need not worry about propagating equalities, and because only a small amount of information has to be shared between the theory solvers in practice – far less, on average, than is passed during the normal execution of the Nelson–Oppen procedure.

A different approach has been proposed by de Moura and Bjørner [59]. These authors also make the equalities part of the model, but instead of letting the SAT solver decide on their values, they attempt to compute a consistent assignment to the theory variables that is as diverse as possible. The equalities are then decided upon by following the assignment to the theory variables.

We mentioned in the aside on p. 253 the option of using BDDs, rather than SAT, for performing lazy encoding. As mentioned, a naive procedure where the predicates label the nodes appeared in [87] and [130]. In the context of hardware verification there have been quite a few publications on multiway decision graphs [53], a generalization of BDDs to various first-order theories.

## Eager Encodings

Some of the algorithms presented in earlier chapters are in fact eager-style decision procedures. The reduction methods for equality logic that are presented in Sect. 4.4 are such algorithms [39, 126]. A similar procedure for difference logic was suggested by Ofer Strichman, Sanjit Seshia, and Randal Bryant in [187]. Procedures that are based on small-domain instantiation (see Sect. 4.5 and a similar procedure for difference logic in [191]) can also be seen as eager encodings, although the connection is less obvious: the encoding is

based not on the skeleton and additional constraints, but rather on an encoding of predicates (equalities, inequalities, etc., depending on the theory) over finite-range variables. The original procedure in [154] used multiterminal BDDs rather than SAT to solve the resulting propositional formula. We should also mention that there are hybrid approaches, combining encodings based on small-domain instantiation and explicit constraints, such as the work by Seshia et al. on difference logic [177].

The first proof-based reduction corresponding to an eager encoding (from integer- and real-valued linear arithmetic) was introduced by Ofer Strichman [186]. The procedure was not presented as part of a more general framework of using deductive rules as described in this chapter. The proof was generated in an eager manner using Fourier–Motzkin variable elimination for the reals and the Omega test for the integers. The example in Sect. 11.3.5 is based on the Boolean Fourier–Motzkin reduction algorithm suggested in [186].

There are only a few publicly available, supported decision procedures based on eager encoding, most notably UCLID [40], which was developed by Randal Bryant, Shuvendu Lahiri, and Sanjit Seshia. As mentioned earlier in this chapter, the eager approach is, at least at the time of writing, considered empirically inferior to the lazy approach.

## 11.6 Glossary

The following symbols were used in this chapter:

Symbol	Refers to ...	First used on page ...
$e(l)$	The propositional encoder of a $\Sigma$ -literal $l$	241
$\alpha(t)$	A truth assignment (either full or partial) to the variables of a formula $t$	241
$lit(\varphi)$	The literals of $\varphi$	244
$lit_i(\varphi)$	Assuming some predefined order on the literals, this denotes the $i$ -th distinct literal in $\varphi$	244
$\alpha$	An assignment (either full or partial) to the literals	244
$Th(lit_i, \alpha)$	See (11.11)	244
$Th(\alpha)$	See (11.12)	244
<i>continued on next page</i>		

<i>continued from previous page</i>		
<b>Symbol</b>	<b>Refers to ...</b>	<b>First used on page ...</b>
$\hat{T}h(\alpha)$	The conjunction over the elements in $Th(\alpha)$	244
$\mathcal{B}$	A Boolean formula. In this chapter, initially set to $e(\varphi)$ , and then strengthened with constraints	245
$t$	For a $\Sigma$ -theory $T$ , $t$ represents a $\Sigma$ -formula (typically a clause) returned by DEDUCTION	245
$\mathcal{B}^i$	The formula $\mathcal{B}$ in the $i$ -th iteration of the loop in Algorithm 11.2.1	246
$P$	A proof – see Definition 11.4	254
$psc(s)$	A proof step constraint – see Definition 11.6. An implication between the antecedents and the consequent of a proof rule	254
$\hat{P}$	See Definition 11.7. A conjunction of $psc(s)$ for all proof steps $s$ in a proof $P$	254
$\Gamma$	An arbitrary inference system	260

## The Satisfiability-Modulo-Theory Library and Standard (SMT-LIB)

The growing interest and need for decision procedures such as those described in this book led to the **SMT-LIB initiative** (short for Satisfiability-Modulo-Theory Library). The main purpose of this initiative was to streamline the research and tool development in the field to which this book is dedicated. For this purpose, the organizers developed the **SMT-LIB standard** [162], which formally specifies the theories that attract enough interest in the research community, and that have a sufficiently large set of publicly available benchmarks. As a second step, the organizers started collecting benchmarks in this format, and today (2008) the SMT-LIB repository includes more than 60 000 benchmarks in the SMT-LIB format, classified into 12 divisions. A third step was to initiate **SMT-COMP**, an annual competition for SMT solvers, with a separate track for each division.

These three steps have promoted the field dramatically: only a few years back, it was very hard to get benchmarks, every tool had its own language standard and hence the benchmarks could not be migrated without translation, and there was no good way to compare tools and methods.<sup>1</sup> These problems have mostly been solved because of the above initiative, and, consequently, the number of tools and research papers dedicated to this field is now steadily growing.

The SMT-LIB initiative was born at FroCoS 2002, the fourth Workshop on Frontiers of Combining Systems, after a proposal by Alessandro Armando. At the time of writing this appendix, it is co-led by Silvio Ranise and Cesare Tinelli, who also wrote the SMT-LIB standard. Clark Barrett, Leonardo de Moura and Cesare Tinelli currently manage the SMT-LIB benchmark repository. The annual SMT-COMP competitions are currently organized by Aaron Stump, Clark Barrett, and Leonardo de Moura.

---

<sup>1</sup> In fact, it was reported in [61] that each tool tended to be the best on its own set of benchmarks.

# B

---

## A C++ Library for Developing Decision Procedures

### B.1 Introduction

A decision procedure is always more than one algorithm. A lot of infrastructure is required to implement even simple decision procedures. We provide a large part of this infrastructure in form of the DPLIB library in order to simplify the development of new procedures. DPLIB is available for download,<sup>1</sup> and consists of the following parts:

- A template class for a basic data structure for graphs, described in Sect. B.2.
- A parser for a simple fragment of first-order logic given in Sect. B.3.
- Code for generating propositional SAT instances in CNF format, shown in Sect. B.4.
- A template for a decision procedure that performs a lazy encoding, described in Sect. B.5.

To begin with, the decision problem (the formula) has to be read as input by the procedure. The way this is done depends on how the decision procedure interfaces with the program that generates the decision problem.

In industrial practice, many decision procedures are embedded into larger programs in the form of a subprocedure. We call programs that make use of a decision procedure as a subprocedure *applications*. If the run time of the decision procedure dominates the total run time of the application, solvers for decision problems are often interfaced to by means of a *file interface*. This chapter provides the basic ingredients for building a decision procedure that uses a file interface. We focus on the C/C++ programming language, as all of the best-performing decision procedures are written in this language.

The components of a decision procedure with a file interface are shown in Fig. B.1. The first step is to *parse* the input file. This means that a sequence of characters is transformed into a *parse tree*. The parse tree is subsequently


---

<sup>1</sup> <http://www.decision-procedures.org/>



checked for type errors (e.g., adding a Boolean to a real number can be considered a type error). This step is called *type checking*. The module of the program that performs the parsing and type-checking phases is usually called the *front end*.

Most of the decision procedures described in this book permit an arbitrary Boolean structure in the formula, and thus have to reason about propositional logic. The best method to do so is to use a modern SAT solver. We explain how to interface to SAT solvers in Sect. B.4. A simple template for a decision procedure that implements an incremental translation to propositional logic, as described in Chap. 11, is given in Sect. B.5.



**Fig. B.1.** Components of a decision procedure that implements a file interface

## B.2 Graphs and Trees

*Graphs* are a basic data structure used by many decision procedures, and can serve as a generalization of many more data structures. As an example, trees and directed acyclic graphs are obvious special cases of graphs. We have provided a template class that implements a generic graph container.

This class has the following design goals:

- It provides a *numbering* of the nodes. Accessing a node by its number is an  $O(1)$  operation. The node numbers are *stable*, i.e., stay the same even if the graph is changed or copied.
- The data structure is optimized for *sparse* graphs, i.e., with few edges. Inserting or removing edges is an  $O(\log k)$  operation, where  $k$  is the number of edges. Similarly, determining if a particular edge exists is also  $O(\log k)$ .
- The nodes are stored densely in a *vector*, i.e., with very little overhead per node. This permits a large number (millions) of nodes. However, adding or removing nodes may invalidate references to already existing nodes.

An instance of a graph named `G` is created as follows:

```
#include "graph.h"
...
graph<graph_nodet<> > G;
```

Initially, the graph is empty. Nodes can be added in two ways: a single node is added using the method `add_node()`. This method adds one node, and returns the number of this node. If a larger number of nodes is to be added, the method `resize(i)` can be used. This changes the number of nodes to  $i$  by either adding or removing an appropriate number of nodes. Means to erase individual nodes are not provided.

The class `graph` can be used for both directed and undirected graphs. Undirected graphs are simply stored as directed graphs where edges always exist in both directions. We write  $a \rightarrow b$  for a directed edge from  $a$  to  $b$ , and  $a \leftrightarrow b$  for an undirected edge between  $a$  and  $b$ .

<b>Class:</b>	<code>graph&lt;T&gt;</code>	
<b>Methods:</b>	<code>add_edge(a, b)</code>	adds $a \rightarrow b$
	<code>remove_edge(a, b)</code>	removes $a \rightarrow b$ , if it exists
	<code>add_undirected</code>	adds $a \leftrightarrow b$
	<code>_edge(a, b)</code>	
	<code>remove_undirected</code>	removes $a \leftrightarrow b$
	<code>_edge(a, b)</code>	
	<code>remove_in_edges(a)</code>	removes $x \rightarrow a$ , for any node $x$
	<code>remove_out_edges(a)</code>	removes $a \rightarrow x$ , for any node $x$
	<code>remove_edges(a)</code>	removes $a \rightarrow x$ and $x \rightarrow a$ , for any node $x$

**Table B.1.** Interface of the template class `graph<T>`

The methods of this template class are shown in Table B.1. The method `has_edge(a, b)` returns `true` if and only if  $a \rightarrow b$  is in the graph. The set of nodes  $x$  such that  $x \rightarrow a$  is returned by `in(a)`, and the set of nodes  $x$  such that  $a \rightarrow x$  is returned by `out(a)`.

The class `graph` provides an implementation of the following two algorithms:

- The set of nodes that are reachable from a given node  $a$  can be computed using the method `visit_reachable(a)`. This method sets the member `.visited` of all nodes that are reachable from node  $a$  to `true`. This member can be set for all nodes to `false` by calling the method `clear_visited()`.
- The shortest path from a given node  $a$  to a node  $b$  can be computed with the method `shortest_path(a, b, p)`, which takes an object  $p$  of type `graph::patht` (a list of node numbers) as its third argument, and stores the shortest path between  $a$  and  $b$  in there. If  $b$  is not reachable from  $a$ , then  $p$  is empty.

### B.2.1 Adding “Payload”

Many algorithms that operate on graphs may need to store additional information per node or per edge. The container class provides a convenient way to do so by defining a new class for this data, and using this new class as a template argument for the template `graph`. As an example, this can be used to define a graph that has an additional string member in each node:

```
#include "graph.h"

class my_nodet {
public:
    std::string name;
};
...

graph<my_nodet> G;
```

Data members can be added to the edges by passing a class type as a second template argument to the template `graph_nodet`. As an example, the following fragment allows a `weight` to be associated with each edge:

```
#include "graph.h"

class my_edget {
    int weight;

    my_edget():weight(0) {
    }
};

class my_nodet {
};
...

graph<my_nodet, my_edget> G;
```

Individual edges can be accessed using the method `edge()`. The following example sets the weight of edge  $a \rightarrow b$  to 10:

```
G.edge(a, b).weight=10;
```

## B.3 Parsing

### B.3.1 A Grammar for First-Order Logic

Many decision problems are stored in a file. The decision procedure is then passed the name of the file. The first step of the program that implements

```

id          : [a-zA-Z_][a-zA-Z0-9_]*+
N-elem    : [0-9]+
Q-elem    : [0-9]*.[0-9]+
infix-function-id : + | - | * | / | mod
boolop-id  : ^ | v | ⇔ | ⇒
infix-relop-id : < | > | ≤ | ≥ | =
quantifier  : ∀ | ∃
term        : id
              | N-elem | Q-elem
              | id ( term-list )
              | term infix-function-id term
              | - term
              | ( term )
formula    : id
              | id ( term-list )
              | term infix-relop-id term
              | quantifier variable-list : formula
              | ( formula )
              | formula boolop-id formula
              | ¬ formula
              | true | false

```

**Fig. B.2.** Simple BNF grammar for formulas

the decision procedure is therefore to parse the file. The file is assumed to follow a particular syntax. We have provided a parser for a simple fragment of first-order logic with quantifiers.

Figure B.2 shows a grammar of this fragment of first-order logic. The grammar in Fig. B.2 uses mathematical notation. The corresponding ASCII representations are listed in Table B.2.

All predicates, variables, and functions have *identifiers*. These identifiers must be declared before they are used. Declarations of variables come with a *type*. These types allow a problem that is in, for example, linear arithmetic over the integers to be distinguished from a problem in linear arithmetic over the reals. Figure B.3 lists the types that are predefined. The *domain*  $\mathbb{U}$  is used for types that do not fit into the other categories.

```

 $\mathbb{B}$   boolean
 $\mathbb{N}_0$  natural
 $\mathbb{Z}$    int
 $\mathbb{R}$    real
 $\mathbb{B}^N$  unsigned [N]
 $\mathbb{B}^N$  signed  [N]
 $\mathbb{U}$    untyped

```

**Fig. B.3.** Supported types and their ASCII representations

Mathematical Symbol	Operation	ASCII
$\neg$	Negation	not, !
$\wedge$	Conjunction	and, &
$\vee$	Disjunction	or,
$\Leftrightarrow$	Biimplication	<=>
$\Rightarrow$	Implication	=>
$<$	Less than	<
$>$	Greater than	>
$\leq$	Less than or equal to	<=
$\geq$	Greater than or equal to	>=
$=$	Equality	=
$\forall$	Universal quantification	forall
$\exists$	Existential quantification	exists
$-$	Unary minus	-
$\cdot$	Multiplication	*
$/$	Division	/
$mod$	Modulo (remainder)	mod
$+$	Addition	+
$-$	Subtraction	-

**Table B.2.** Built-in function symbols

Table B.2 also defines the precedence of the built-in operators: the operators with higher precedence are listed first, and the precedence levels are separated by horizontal lines. All operators are left-associative.

### B.3.2 The Problem File Format

The input files for the parser consist of a sequence of *declarations* (Fig. B.4 shows an example). All variables, functions, and predicates are declared. The declarations are separated by semicolons, and the elements in each declaration are separated by commas. Each variable declaration is followed by a type (as listed in Fig. B.3), which specifies the type of all variables in that declaration.

A declaration may also define a formula. Formulas are *named* and *tagged*. Each entry starts with the name of the formula, followed by a colon and one of the keywords `theorem`, `axiom`, or `formula`. The keyword is followed by a formula. Note that the formulas are not necessarily *closed*: the formula `simplex` contains the unquantified variables `i` and `j`. Variables that are not quantified explicitly are implicitly quantified with a universal quantifier.

```

a, b, x, p, n: int;
el: natural;
pi: real;
i, j: real;
u: untyped;          -- an untyped variable
abs: function;
prime, divides: predicate;

absolute: axiom      forall a: ((a >=0 ==> abs(a) = a) and
                               (a < 0 ==> abs(a) = -a)) ==>
                               (exists el: el = abs(a));
divides: axiom       (forall a, b: divides (a, b) <=>
                       exists x: b = a * x);
simplex: formula      (i + 5*j <= 3) and
                       (3*i < 3.7) and
                       (i > -1) and (j > 0.12)

```

**Fig. B.4.** A realistic example

### B.3.3 A Class for Storing Identifiers

Decision problems often contain a large set of variables, which are represented by identifier strings. The main operation on these identifiers is comparison. We therefore provide a specialized string class that features string comparison in time  $O(1)$ . This is implemented by storing all identifiers inside a hash table. Comparing strings then reduces to comparing indices for that table.

Identifiers are stored in objects of type `dstring`. This class offers most of the methods that the other string container classes feature, with the exception of any method that modifies the string. Instances of type `dstring` can be copied, compared, ordered, and destroyed in time  $O(1)$ , and use as much space as an integer variable.

### B.3.4 The Parse Tree

The parse tree is stored in a graph class `ast::astt` and is generated from a file as follows (Fig. B.5):

1. Create an instance of the class `ast::astt`.
2. Call the method `parse(file)` with the name of the file as an argument.

The method returns `true` if an error was encountered during parsing.

The class `ast::astt` is a specialized form of a graph, and stores nodes of type `ast::nodet`. The root node is returned by the method `root()` of the class `ast::astt`. Each node stores the following information:

```

#include "parsing/ast.h"

...

ast::astt ast;

if(ast.parse(argv[1])) {
    std::cerr << "parsing_failed" << std::endl;
    exit(1);
}

```

**Fig. B.5.** Generating a parse tree

1. Each node has a numeric label (an integer). This is used to distinguish the operators and the terminal symbols. Table B.3 contains a list of the symbolic constants that are used for the numeric labels.
2. Nodes that contain identifiers or a numeric constant also have a string label, which is of type `dstring` (see Sect. B.3.3). We use strings for the numeric constants instead of the numeric types offered by C++ in order to support unbounded numbers.
3. Each node may have up to two child nodes.

As described in Sect. B.2, the nodes of the graph are numbered. In fact, the `ast::nodet` class is only a wrapper around these numbers, and thus can be copied efficiently. The methods it offers are shown in Table B.4. The methods `c1()` and `c2()` return `NIL` if there is no first or second child node, respectively.

For convenience, the `ast::astt` class provides a *symbol table*, which is a mapping from the set of identifiers to their types. Given an identifier  $s$ , the method `get_type_node(s)` returns the node in the parse tree that corresponds to the type of  $s$ .

## B.4 CNF and SAT

### B.4.1 Generating CNF

The library provides algorithms for converting propositional logic into CNF using Tseitin's method (see Sect. 1.3). The resulting clauses can be passed directly to a propositional SAT solver. Alternatively, they can be written to disk in the DIMACS format. The interface to both back ends is defined in the `propt` base class. This class is used wherever the specific propositional back end is to be left unspecified. Literals (i.e., variables or their negations) are

Name	Used for
N_IDENTIFIER	Identifier
N_INTEGER	Integer constant
N_RATIONAL	Rational constant
N_INT	Integer type
N_REAL	Real type
N_BOOLEAN	Boolean type
N_UNSIGNED	Unsigned type
N_SIGNED	Signed type
N_AXIOM	Axiom
N_DECLARATION	Declaration
N_THEOREM	Theorem
N_CONJUNCTION	$\wedge$
N_DISJUNCTION	$\vee$
N_NEGATION	$\neg$
N_BIIMPLICATION	$\iff$
N_IMPLICATION	$\implies$
N_TRUE	True
N_FALSE	False
N_ADDITION	$+$
N_SUBTRACTION	$-$
N_MULTIPLICATION	$*$
N_DIVISION	$/$
N_MODULO	mod
N_UMINUS	Unary minus
N_LOWER	$<$
N_GREATER	$>$
N_LOWEREQUAL	$\leq$
N_GREATEREQUAL	$\geq$
N_EQUAL	$=$
N_FORALL	$\forall$
N_EXISTS	$\exists$
N_LIST	A list of nodes
N_PREDICATE	Predicate
N_FUNCTION	Function

**Table B.3.** Numeric labels of nodes and their meanings



<b>Class:</b>	ast::nodet	
<b>Methods:</b>	id()	Returns the numeric label
	string()	Returns the string label
	c1()	Returns the first child node
	c2()	Returns the second child node
	number()	Returns the number of the node
	is_nil()	Returns true if the node is NIL

**Table B.4.** Interface of the class ast::nodet

stored in objects of type literal\_t. The constants TRUE and FALSE are returned by const\_literal(true) and const\_literal(false), respectively.

<b>Class:</b>	propt	
<b>Methods:</b>	land( $a, b$ )	Returns a literal $l$ with $l \iff a \wedge b$
	land( $v$ )	Given a vector $v = \langle v_1, \dots, v_n \rangle$ , returns a literal $l$ with $l \iff \bigwedge_i v_i$
	lor( $a, b$ )	Returns a literal $l$ with $l \iff a \vee b$
	lor( $v$ )	Given a vector $v = \langle v_1, \dots, v_n \rangle$ , returns a literal $l$ with $l \iff \bigvee_i v_i$
	lxor( $a, b$ )	Returns a literal $l$ with $l \iff a \oplus b$
	lnot( $a, b$ )	Returns a literal $l$ with $l \iff \neg a$
	lnand( $a, b$ )	Returns a literal $l$ with $l \iff \neg(a \wedge b)$
	lnor( $a, b$ )	Returns a literal $l$ with $l \iff \neg(a \vee b)$
	lequal( $a, b$ )	Returns a literal $l$ with $l \iff (a \iff b)$
	limplies( $a, b$ )	Returns a literal $l$ with $l \iff (a \implies b)$
	lselect( $a, b, c$ )	Returns a literal $l$ with $(a \implies (l \iff b)) \wedge (\neg a \implies (l \iff c))$
	set_equal( $a, b$ )	Adds the constraint $a \iff b$
	new_variable()	Returns a new variable
	const_literal( $c$ )	Returns a literal with a constant Boolean truth value given by $c$

**Table B.5.** Interface of the class propt

The interface of the class propt is specified in Table B.5. The classes satcheck\_t and dimacs\_cnft are derived from this class. An implementation of a state-of-the-art propositional SAT solver is given by the class satcheck\_t. The additional methods it provides are shown in Table B.6. The class dimacs\_cnft is used to store the clauses and dump them into a text file that uses the DIMACS CNF format. Its interface is given in Table B.7.

<b>Class:</b>	satcheckt, derived from propt	
<b>Methods:</b>	prop_solve()	Returns P-SATISFIABLE if the formula is SAT
	l_get( <i>l</i> )	Returns the value of <i>l</i> in the satisfying assignment
	solver_text()	Returns a string that identifies the solver

**Table B.6.** Interface of the class satcheckt

<b>Class:</b>	dimacs_cnft, derived from propt	
<b>Methods:</b>	write_dimacs_cnft( <i>s</i> )	Dumps the formula in DIMACS CNF format into the stream <i>s</i>

**Table B.7.** Interface of the class dimacs\_cnft

## B.4.2 Converting the Propositional Skeleton

The **propositional skeleton** (see Chap. 11) of a parse tree can be generated using the class `skeletont`. This offers an operator `()`, which can be applied as follows, where `root_node` is the root node of a formula, and `prop` is an instance of `propt`:

```
#include "sat/skeleton.h"

...

skeletont skeleton;

skeleton(root_node, prop);
```

Besides converting the propositional part, the method also generates a vector `skeleton.nodes`, where each element corresponds to a node in the parse tree. Each node has two attributes:

- The attribute `type` is one of `PROPOSITIONAL` or `THEORY`, and distinguishes the skeleton from the theory atoms.
- In the case of a skeleton node, the attribute `l` is the literal that encodes the node.

## B.5 A Template for a Lazy Decision Procedure

The library provides two templates for decision procedures that compute a propositional encoding of a given formula  $\varphi$  in the lazy manner. These

algorithms are described in detail under the names LAZY-DPLL (Algorithm 11.2.2) and DPLL( $T$ ) (Algorithm 11.2.3) in Chap. 11.

We first define a common interface for any kind of decision procedure. This interface is defined by a class `decision_procedure_t` (Table B.8). This class offers a method `is_satisfiable( $\varphi$ )`, which returns `TRUE` if and only if the formula  $\varphi$  is satisfiable. If so, one may call the methods `print_assignment( $s$ )` and `get( $n$ )`. The method `print_assignment( $s$ )` dumps the entire satisfying assignment into a stream, whereas `get( $n$ )` permits querying the value of an individual node  $n$  of  $\varphi$ .

<b>Class:</b>	<code>decision_procedure_t</code>
<b>Methods:</b>	<code>is_satisfiable(<math>\varphi</math>)</code> Returns <code>TRUE</code> if the formula $\varphi$ is found to be SAT
	<code>print_assignment(<math>s</math>)</code> Dumps the satisfying assignment into the stream $s$
	<code>get(<math>n</math>)</code> Returns the value assigned to node $n$ of $\varphi$

**Table B.8.** Interface of the class `decision_procedure_t`

<b>Class:</b>	<code>lazy_dp11t</code> , derived from <code>decision_procedure_t</code>
<b>Methods:</b>	<code>assignment(<math>n, v</math>)</code> This method is called by the SAT solver for every assignment to a $\Sigma$ -literal in $\varphi$ . The node it corresponds to is $n$ ; the value assigned is given by $v$ .
	<code>deduce()</code> This method is called once a satisfying assignment to the current propositional encoding is found.
	<code>add_clause(<math>c</math>)</code> Called by <code>deduce()</code> to add a clause as a consequence of a $T$ -inconsistent assignment
<b>Members:</b>	<code>f</code> A copy of $\varphi$
	<code>skeleton</code> An instance of <code>skeleton_t</code>

**Table B.9.** Interface of the classes `lazy_dp11t` and `dp11.tt`, which are implementations of LAZY-DPLL (Algorithm 11.2.2) and DPLL( $T$ ) (Algorithm 11.2.3). The theory  $T$  is assumed to be defined over a signature  $\Sigma$

The templates that we have provided implement two of the algorithms given in Chap. 11: LAZY-DPLL and DPLL( $T$ ). These templates include the conversion of the propositional skeleton of  $\varphi$  into CNF, and the interface to

the SAT solver. We provide a common interface to both algorithms, which is given in Table B.9.

<b>Class:</b>	<code>dpll_tt</code> , derived from <code>decision_procedure_t</code>	
<b>Methods:</b>	<code>deduce()</code>	This method is called by the SAT solver to check a partial assignment for $T$ -consistency.
	<code>add_clause(c)</code>	Called to add a clause as consequence of assignment
	<code>theory_implication(n, v)</code>	Called to communicate a $T$ -implication to the SAT solver: $n$ is the node implied, and $v$ is the value.
<b>Members:</b>	<code>f</code>	A copy of $\varphi$
	<code>skeleton</code>	An instance of <code>skeleton_t</code>

**Table B.10.** Interface of the class `dpll_tt`, an implementation of  $DPLL(T)$  (Algorithm 11.2.3)

The only part that is left open is the interface to the decision procedure for the conjunction of  $\Sigma$ -literals. In the case of both algorithms, this is the method `deduce()`. The assignment to the  $\Sigma$ -literals is passed from the SAT solver to the deductive engine by means of calls to the method `assignment(n, v)`, where  $n$  is the node and  $v$  is the value that is assigned.

The method `deduce()` inspects this assignment to the  $\Sigma$ -literals. If it is found to be consistent, `deduce()` is expected to return `TRUE`. Otherwise, it is expected to add appropriate constraints using the method `add_clause`, and to return `FALSE`.

In the case of `LAZY-DPLL`, `deduce()` is called only for full assignments, whereas  $DPLL(T)$  may call `deduce()` for partial assignments.

---

## References

1. W. Ackermann. *Solvable cases of the Decision Problem*. Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam, 1954.
2. V. S. Adve and J. M. Mellor-Crummey. Using integer sets for data-parallel program analysis and optimization. In *PLDI*, pages 186–198, 1998.
3. S. Akers. Binary decision diagrams. *IEEE Transactions on Computers*, C-27(6):509–516, 1978.
4. A. Armando, C. Castellini, and E. Giunchiglia. SAT-based procedures for temporal reasoning. In *Proceedings of the 5-th European Conference on Planning (ECP)*, volume 1809 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 1999.
5. A. Armando and E. Giunchiglia. Embedding complex decision procedures inside an interactive theorem prover. *Annals of Mathematics and Artificial Intelligence*, 8(3–4):475–502, 1993.
6. G. Audemard, P. Bertoli, A. Cimatti, A. Kornilowicz, and R. Sebastiani. A SAT based approach for solving formulas over Boolean and linear mathematical propositions. In *Proceedings of the 18-th International Conference on Automated Deduction (CADE)*, volume 2392 of *Lecture Notes in Computer Science*, pages 195–210. Springer, 2002.
7. D. Babic and F. Hutter. Spear. Proceedings of the SAT 2007 competition, Jan 2007.
8. L. Bachmair and A. Tiwari. Abstract congruence closure and specializations. In *CADE 2000*, volume 1831 of *Lecture Notes in Computer Science*, pages 64–78. Springer, 2000.
9. C. Barrett and S. Berezin. CVC Lite: A new implementation of the cooperating validity checker. In *CAV 04: International Conference on Computer-Aided Verification*, volume 3114 of *Lecture Notes in Computer Science*. Springer, 2004.
10. C. Barrett, D. Dill, and J. Levitt. Validity checking for combinations of theories with equality. In *Proceedings FMCAD 1996*, volume 1166 of *Lecture Notes in Computer Science*. Springer, 1996.
11. C. Barrett and C. Tinelli. CVC3. In *Proceedings of the 19-th International Conference on Computer Aided Verification (CAV)*, volume 4590 of *Lecture Notes in Computer Science*, pages 298–302. Springer, 2007.

12. C. W. Barrett, D. L. Dill, and J. R. Levitt. A decision procedure for bit-vector arithmetic. In *Design Automation Conference (DAC)*, pages 522–527. ACM Press, 1998.
13. C. W. Barrett, D. L. Dill, and A. Stump. Checking satisfiability of first-order formulas by incremental translation to SAT. In *Proceedings of the 14-th International Conference on Computer Aided Verification (CAV)*, volume 2404 of *Lecture Notes in Computer Science*, pages 236–249. Springer, 2002.
14. C. W. Barrett, D. L. Dill, and A. Stump. A generalization of Shostak’s method for combining decision procedures. In *FroCos*, pages 132–146, 2002.
15. R. Bayardo and R. Schrag. Using CSP look-back techniques to solve real-world SAT instances. In *14-th National (US) Conference on Artificial Intelligence (AAAI)*, pages 203–208, 1997.
16. P. Beame, H. Kautz, and A. Sabharwal. Towards understanding and harnessing the potential of clause learning. *Journal of Artificial Intelligence Research*, 22:319–351, 2004.
17. M. Benedikt, T. W. Reps, and S. Sagiv. A decidable logic for describing linked data structures. In *Programming Languages and Systems, 8-th European Symposium on Programming, ESOP’99*, volume 1576 of *Lecture Notes in Computer Science*, pages 2–19. Springer, 1999.
18. J. Berdine, C. Calcagno, and P. W. O’Hearn. A decidable fragment of separation logic. In *FSTTCS 2004: Foundations of Software Technology and Theoretical Computer Science, 24-th International Conference*, volume 3328 of *Lecture Notes in Computer Science*, pages 97–109. Springer, 2004.
19. J. Berdine, C. Calcagno, and P. W. O’Hearn. Symbolic execution with separation logic. In *Programming Languages and Systems, 3-rd Asian Symposium, (APLAS)*, volume 3780 of *Lecture Notes in Computer Science*, pages 52–68. Springer, 2005.
20. J. Berdine, C. Calcagno, and P. W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In *Formal Methods for Components and Objects (FMCO)*, volume 4111 of *Lecture Notes in Computer Science*, pages 115–137. Springer, 2006.
21. D. Beyer, T. A. Henzinger, and G. Théoduloz. Lazy shape analysis. In *Computer Aided Verification, 18-th International Conference, CAV 2006, Seattle, WA, USA, August 17-20, 2006, Proceedings*, volume 4144 of *Lecture Notes in Computer Science*, pages 532–546. Springer, 2006.
22. A. Biere. Resolve and expand. In *Proceedings of the 7-th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, 2004.
23. A. Biere, A. Cimatti, E. Clarke, and Y. Zhu. Symbolic model checking without BDDs. In *Proceedings of the Workshop on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1579 of *Lecture Notes in Computer Science*, pages 193–207. Springer, 1999.
24. J. D. Bingham and Z. Rakamaric. A logic and decision procedure for predicate abstraction of heap-manipulating programs. In *Verification, Model Checking, and Abstract Interpretation (VMCAI)*, volume 3855 of *Lecture Notes in Computer Science*, pages 207–221. Springer, 2006.
25. C. Binnig, D. Kossmann, and E. Lo. Reverse query processing. In *Proceedings of the 23-rd International Conference on Data Engineering (ICDE)*, pages 506–515. IEEE, 2007.

26. C. Binnig, D. Kossmann, E. Lo, and M. T. Özsu. QAGen: generating query-aware test databases. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 341–352. ACM, 2007.
27. M. P. Bonacina, S. Ghilardi, E. Nicolini, S. Ranise, and D. Zucchelli. Decidability and undecidability results for Nelson–Oppen and rewrite-based decision procedures. In *Automated Reasoning (IJCAR)*, volume 4130 of *Lecture Notes in Computer Science*, pages 513–527. Springer, 2006.
28. M. Bozzano, R. Bruttomesso, A. Cimatti, T. Junttila, S. Ranise, P. van Rossum, and R. Sebastiani. Efficient satisfiability modulo theories via delayed theory combination. In *Proceedings of the 17-th International Conference on Computer Aided Verification (CAV)*, volume 3576 of *Lecture Notes in Computer Science*, pages 335–349, 2005.
29. K. S. Brace, R. L. Rudell, and R. E. Bryant. Efficient implementation of a BDD package. In *DAC*, pages 40–45, 1990.
30. A. R. Bradley and Z. Manna. *The Calculus of Computation*. Springer, 2007.
31. A. R. Bradley, Z. Manna, and H. B. Sipma. What’s decidable about arrays? In *Verification, Model Checking, and Abstract Interpretation, 7-th International Conference (VMCAI)*, volume 3855 of *Lecture Notes in Computer Science*, pages 427–442. Springer, 2006.
32. R. Brinkmann and R. Drechsler. RTL-datapath verification using integer linear programming. In *Proceedings of VLSI Design*, pages 741–746. IEEE, 2002.
33. R. Brummayer and A. Biere. C32SAT: Checking C expressions. In *Computer Aided Verification (CAV)*, volume 4590 of *Lecture Notes in Computer Science*, pages 294–297. Springer, 2007.
34. R. Bruttomesso, A. Cimatti, A. Franzén, A. Griggio, Z. Hanna, A. Nadel, A. Palti, and R. Sebastiani. A lazy and layered SMT(BV) solver for hard industrial verification problems. In *Computer Aided Verification (CAV)*, volume 4590 of *Lecture Notes in Computer Science*, pages 547–560. Springer, 2007.
35. R. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Transactions on Computers*, C-35(12):1035–1044, 1986.
36. R. Bryant. Symbolic Boolean manipulation with ordered binary decision diagrams. *ACM Computing Surveys*, 24(3):293–318, 1992.
37. R. Bryant and Y. Chen. Verification of arithmetic circuits with binary moment diagrams. In *32-nd ACM/IEEE Design Automation Conference*, June 1995.
38. R. Bryant, S. German, and M. Velev. Exploiting positive equality in a logic of equality with uninterpreted functions. In *Proceedings of the 11-th International Conference on Computer Aided Verification (CAV)*, volume 1633 of *Lecture Notes in Computer Science*. Springer, 1999.
39. R. Bryant and M. Velev. Boolean satisfiability with transitivity constraints. In *Proceedings of the 12-th International Conference on Computer Aided Verification (CAV)*, volume 1855 of *Lecture Notes in Computer Science*. Springer, 2000.
40. R. E. Bryant, S. K. Lahiri, and S. A. Seshia. Modeling and verifying systems using a logic of counter arithmetic with lambda expressions and uninterpreted functions. In *Proceedings of the 14-th International Conference on Computer Aided Verification (CAV)*, volume 2404 of *Lecture Notes in Computer Science*. Springer, 2002.
41. H. K. Büning, M. Karpinski, and A. Flögel. Resolution for quantified Boolean formulas. *Information and Computation*, 117(1):12–18, 1995.

42. R. Burstall. Some techniques for proving correctness of programs which alter data structures. *Machine Intelligence*, 7:23–50, 1972.
43. C. Calcagno, D. Distefano, P. W. O’Hearn, and H. Yang. Beyond reachability: Shape abstraction in the presence of pointer arithmetic. In *Static Analysis, 13-th International Symposium (SAS)*, volume 4134 of *Lecture Notes in Computer Science*, pages 182–203. Springer, 2006.
44. C. Calcagno, H. Yang, and P. W. O’Hearn. Computability and complexity results for a spatial assertion language for data structures. In *Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 2245 of *Lecture Notes in Computer Science*, pages 108–119. Springer, 2001.
45. P. Cheeseman, B. Kanefsky, and W. M. Taylor. Where the really hard problems are. In *Proceedings of the 12-th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 331–337, 1991.
46. E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, Cambridge, MA, 1999.
47. E. Clarke, K. McMillan, X. Zhao, M. Fujita, and J. Yang. Spectral transforms for large Boolean functions with applications to technology mapping. In *Proceedings of the 30-th Design Automation Conference (DAC93)*, pages 54–60. ACM Press, 1993.
48. E. M. Clarke, D. Kroening, and F. Lerda. A tool for checking ANSI-C programs. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 2988 of *Lecture Notes in Computer Science*, pages 168–176. Springer, 2004.
49. B. Cook, D. Kroening, and N. Sharygina. Accurate theorem proving for program verification. In *Proceedings of ISoLA 2004*, volume 4313 of *Lecture Notes in Computer Science*, pages 96–114. Springer, 2006.
50. S. Cook. The complexity of theorem-proving procedures. In *Proceedings of the 3-rd Annual ACM Symposium on Theory of Computing*, pages 151–158, 1971.
51. D. C. Cooper. Theorem proving in arithmetic without multiplication. In *Machine Intelligence*, pages 91–100. Edinburgh University Press, 1972.
52. F. Copty, L. Fix, R. Fraer, E. Giunchiglia, G. Kamhi, A. Tacchella, and M. Y. Vardi. Benefits of bounded model checking at an industrial setting. In *Proceedings of the 13-th International Conference on Computer Aided Verification (CAV)*, volume 2102 of *Lecture Notes in Computer Science*, pages 436–453, 2001.
53. F. Corella, Z. Zhou, X. Song, M. Langevin, and E. Cerny. Multiway decision graphs for automated hardware verification. *Formal Methods in System Design*, 10(1):7–46, 1997.
54. T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*, chapter 25.3, pages 532–536. MIT Press, 2000.
55. G. B. Danzig. *Linear Programming and Extensions*. Princeton University Press, 1963.
56. M. Davis, G. Logemann, and D. Loveland. A machine program for theorem-proving. *Communications of the ACM*, 5:394–397, 1962.
57. M. Davis and H. Putnam. A computing procedure for quantification theory. *Journal of the ACM*, 7:201–215, 1960.
58. L. de Moura. System description: Yices 0.1. Technical report, Computer Science Laboratory, SRI International, 2005.
59. L. de Moura and N. Bjørner. Model-based theory combination. In *Satisfiability Modulo Theories (SMT)*, 2007.



60. L. de Moura and H. Ruess. Lemmas on demand for satisfiability solvers. In *Proceedings of SAT'02*, 2002.
61. L. de Moura and H. Ruess. An experimental evaluation of ground decision procedures. In *Proceedings of the 16-th International Conference on Computer Aided Verification (CAV)*, volume 3114 of *Lecture Notes in Computer Science*, pages 162–174. Springer, 2004.
62. L. de Moura, H. Rueß, and N. Shankar. Justifying equality. In *Proceedings of the Second Workshop on Pragmatics of Decision Procedures in Automated Reasoning (PDPAR)*, 2004.
63. L. M. de Moura and N. Bjørner. Efficient E-matching for SMT solvers. In *Automated Deduction – CADE-21*, volume 4603 of *Lecture Notes in Computer Science*, pages 183–198. Springer, 2007.
64. N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In *Handbook of Theoretical Computer Science (Volume B): Formal Models and Semantics*, pages 243–320. MIT Press, 1990.
65. D. Detlefs, G. Nelson, and J. B. Saxe. Simplify: a theorem prover for program checking. *Journal of the ACM*, 52(3):365–473, May 2005.
66. A. Deutsch. Interprocedural may-alias analysis for pointers: Beyond  $k$ -limiting. In *PLDI*, pages 230–241, 1994.
67. D. Distefano, P. W. O’Hearn, and H. Yang. A local shape analysis based on separation logic. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 3920 of *Lecture Notes in Computer Science*, pages 287–302. Springer, 2006.
68. A. Dolzmann, T. Sturm, and V. Weispfenning. Real quantifier elimination in practice. Technical Report MIP9720, FMI, Universität Passau, Dec. 1997.
69. P. J. Downey, R. Sethi, and R. E. Tarjan. Variations on the common subexpression problem. *J. ACM*, 27(4):758–771, October 1980.
70. B. Dutertre and L. de Moura. A fast linear-arithmetic solver for DPLL(T). In *Proceedings of the 18-th International Conference on Computer Aided Verification (CAV)*, volume 4144 of *Lecture Notes in Computer Science*, pages 81–94. Springer, 2006.
71. B. Dutertre and L. de Moura. Integrating simplex with DPLL(T). Technical Report SRI-CSL-06-01, Stanford Research Institute (SRI), 2006.
72. G. C. E. Balas, S. Ceria and N. Natraj. Gomory cuts revisited. *Operations Research Letters*, 19:1–9, 1996.
73. N. Eén and N. Sörensson. An extensible SAT-solver [ver 1.2]. In *Proceedings of Theory and Applications of Satisfiability Testing*, volume 2919 of *Lecture Notes in Computer Science*, pages 512–518. Springer, 2003.
74. J. Filliatre, S. Owre, H. Ruess, and N. Shankar. ICS: Integrated canonizer and solver. In *Proceedings of the 13-th International Conference on Computer Aided Verification (CAV)*, Lecture Notes in Computer Science. Springer, 2001.
75. M. J. Fischer and M. O. Rabin. Super-exponential complexity of Presburger arithmetic. In *Proceedings of the SIAM-AMS Symposium in Applied Mathematics*, volume 7, pages 27–41, 1974.
76. C. Flanagan, R. Joshi, X. Ou, and J. B. Saxe. Theorem proving using lazy proof explication. In *Computer Aided Verification (CAV)*, volume 2725 of *Lecture Notes in Computer Science*, pages 355–367. Springer, 2003.
77. V. Ganesh, S. Berezin, and D. Dill. Deciding Presburger arithmetic by model checking and comparisons with other methods. In *Formal Methods in*

- Computer-Aided Design (FMCAD)*, volume 2517 of *Lecture Notes in Computer Science*, pages 171–186. Springer, 2002.
78. V. Ganesh and D. L. Dill. A decision procedure for bit-vectors and arrays. In *Computer Aided Verification (CAV)*, volume 4590 of *Lecture Notes in Computer Science*, pages 519–531. Springer, 2007.
  79. H. Ganzinger, G. Hagen, R. Nieuwenhuis, A. Oliveras, and C. Tinelli. DPLL(T): Fast decision procedures. In *Proceedings of the 16-th International Conference on Computer Aided Verification (CAV)*, volume 3114 of *Lecture Notes in Computer Science*, pages 175–188. Springer, 2004.
  80. A. V. Gelder and Y. Tsuji. Incomplete thoughts about incomplete satisfiability procedures. In *2-nd DIMACS Challenge Workshop: Cliques, Coloring and Satisfiability*, 1993.
  81. R. Gershman, M. Koifman, and O. Strichman. Deriving small unsatisfiable cores with dominators. In *Proceedings of the 18-th International Conference on Computer Aided Verification (CAV)*, volume 4144 of *Lecture Notes in Computer Science*, pages 109–122. Springer, 2006.
  82. R. Gershman and O. Strichman. HaifaSat: A new robust SAT solver. In *1-st International Haifa Verification Conference*, volume 3875 of *Lecture Notes in Computer Science*, pages 76–89. Springer, 2005.
  83. M. Ghasemzadeh, V. Klotz, and C. Meinel. Embedding memoization to the semantic tree search for deciding QBFs. In *Australian Conference on Artificial Intelligence*, pages 681–693, 2004.
  84. S. Ghilardi, E. Nicolini, S. Ranise, and D. Zucchelli. Deciding extensions of the theory of arrays by integrating decision procedures and instantiation strategies. In *Logics in Artificial Intelligence (JELIA)*, volume 4160 of *Lecture Notes in Computer Science*, pages 177–189. Springer, 2006.
  85. E. Giunchiglia, M. Narizzano, and A. Tacchella. Quantified Boolean Formulas satisfiability library (QBFLIB), 2001. [www.qbflib.org](http://www.qbflib.org).
  86. F. Giunchiglia and R. Sebastiani. Building decision procedures for modal logics from propositional decision procedures – the case study of modal K. In *CADE-13*, volume 1104 of *Lecture Notes in Computer Science*, New Brunswick, NJ, USA, August 1996. Springer.
  87. A. Goel, K. Sajid, H. Zhou, A. Aziz, and V. Singhal. BDD based procedures for a theory of equality with uninterpreted functions. In *CAV98*, volume 1427 of *Lecture Notes in Computer Science*. Springer, 1998.
  88. E. Goldberg and Y. Novikov. Berkmin: A fast and robust SAT-solver. In *Design, Automation and Test in Europe Conference and Exhibition (DATE)*, page 142, 2002.
  89. R. Gomory. An algorithm for integer solutions to linear problems. In *Recent Advances in Mathematical Programming*, pages 269–302, New-York, 1963. McGraw-Hill.
  90. A. Haken. The intractability of resolution. *Theoretical Computer Science*, 39:297–308, 1985.
  91. J. Harrison. *Handbook of Practical Logic and Automated Reasoning*. Cambridge University Press, 2008.
  92. F. Hillier and G. Lieberman. *Introduction to Mathematical Programming*. McGraw-Hill, 1990.
  93. C. A. R. Hoare and N. Wirth. An axiomatic definition of the programming language PASCAL. *Acta Informatica*, 2(4):335–355, December 1973.

94. R. Hojati, A. Isles, D. Kirkpatrick, and R. Brayton. Verification using uninterpreted functions and finite instantiations. In *Proceedings of the 1-st International Conference on Formal Methods in Computer-Aided Design (FMCAD)*, volume 1166 of *Lecture Notes in Computer Science*, pages 218 – 232. Springer, 1996.
95. R. Hojati, A. Kuehlmann, S. German, and R. Brayton. Validity checking in the theory of equality using finite instantiations. In *Proceedings of the International Workshop on Logic Synthesis*, 1997.
96. J. N. Hooker. Solving the incremental satisfiability problem. *Journal of Logic Programming*, 15:177–186, 1993.
97. I. Horrocks. The FaCT system. In H. de Swart, editor, *Proceedings of the TABLEAUX-98*, volume 1397 of *Lecture Notes in Artificial Intelligence*, pages 307–312. Springer, 1998.
98. J. Huang. MUP: A minimal unsatisfiability prover. In *Proceedings of the 10-th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pages 432–437, 2005.
99. G. Huet and D. Oppen. Equations and rewrite rules: A survey. In *Formal Language Theory: Perspectives and Open Problems*, pages 349–405. Academic Press, 1980.
100. A. P. Hurst, P. Chong, and A. Kuehlmann. Physical placement driven by sequential timing analysis. In *International Conference on Computer-Aided Design (ICCAD)*, pages 379–386. IEEE Computer Society/ACM, 2004.
101. N. Immerman, A. M. Rabinovich, T. W. Reps, S. Sagiv, and G. Yorsh. The boundary between decidability and undecidability for transitive-closure logics. In *Computer Science Logic (CSL)*, volume 3210 of *Lecture Notes in Computer Science*, pages 160–174. Springer, 2004.
102. G. Johnson. Separating the insolvable and the merely difficult. *New York Times*, July 13 1999.
103. T. Jussila, A. Biere, C. Sinz, D. Kroening, and C. M. Wintersteiger. A first step towards a unified proof checker for QBF. In *Theory and Applications of Satisfiability Testing (SAT)*, volume 4501 of *Lecture Notes in Computer Science*, pages 201–214. Springer, 2007.
104. M. Kaufmann, P. Manolios, and J. S. Moore. *Computer-Aided Reasoning: An Approach*. Kluwer Academic, 2000.
105. J. Kim, J. Whitemore, J. Silva, and K. Sakallah. Incremental Boolean satisfiability and its application to delay fault testing. In *IEEE/ACM International Workshop on Logic Synthesis (IWLS)*, June 1999.
106. S. Kirkpatrick and B. Selman. Critical behavior in the satisfiability of random Boolean expressions. *Science*, 264(5163):1297–1301, 1994.
107. D. Kroening, E. Clarke, and K. Yorav. Behavioral consistency of C and Verilog programs using bounded model checking. In *Proceedings of DAC 2003*, pages 368–371. ACM Press, 2003.
108. S. Krstić and S. Conchon. Canonization for disjoint unions of theories. *Information and Computation*, 199:87–106, 2005.
109. S. Krstić and A. Goel. Architecting solvers for SAT modulo theories: Nelson–Oppen with DPLL. In *Proceedings Frontiers of Combining Systems, 6-th International Symposium, FroCoS 2007*, volume 4720 of *Lecture Notes in Artificial Intelligence*, pages 1–27. Springer, 2007.

110. V. Kuncak and M. C. Rinard. Existential heap abstraction entailment is undecidable. In *Static Analysis (SAS)*, volume 2694 of *Lecture Notes in Computer Science*, pages 418–438. Springer, 2003.
111. R. Kurshan. *Computer Aided Verification of Coordinating Processes*. Princeton University Press, 1994.
112. S. K. Lahiri, R. E. Bryant, A. Goel, and M. Talupur. Revisiting positive equality. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 2988 of *Lecture Notes in Computer Science*, pages 1–15. Springer, 2004.
113. S. K. Lahiri and S. Qadeer. Verifying properties of well-founded linked lists. In *Principles of Programming Languages (POPL)*, pages 115–126. ACM, 2006.
114. S. K. Lahiri and S. Qadeer. Back to the future: Revisiting precise program verification using SMT solvers. In *Principles of Programming Languages (POPL)*. ACM, 2008.
115. C. Y. Lee. Representation of switching circuits by binary-decision programs. *Bell Systems Technical Journal*, 38:985–999, 1959.
116. S. Lee and D. Plaisted. Eliminating duplication with the hyper-linking strategy. *Journal of Automated Reasoning*, 9:25–42, 1992.
117. K. R. M. Leino. *Toward reliable modular programs*. PhD thesis, CalTech, 1995. Available as Technical Report Caltech-CS-TR-95-03.
118. I. Lynce and J. Marques-Silva. On computing minimum unsatisfiable cores. In *Proceedings of the International Symposium on Theory and Applications of Satisfiability Testing (SAT)*, pages 305–310, 2004.
119. M. Mahfoudh. *On Satisfiability Checking for Difference Logic*. PhD thesis, Verimag, France, 2003.
120. M. Mahfoudh, P. Niebert, E. Asarin, and O. Maler. A Satisfiability Checker for Difference Logic. In *Proceedings of 5-th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 222–230, 2002.
121. R. Manevich, E. Yahav, G. Ramalingam, and S. Sagiv. Predicate abstraction and canonical abstraction for singly-linked lists. In *Verification, Model Checking, and Abstract Interpretation, 6-th International Conference (VMCAI)*, volume 3385 of *Lecture Notes in Computer Science*, pages 181–198. Springer, 2005.
122. P. Manolios, S. K. Srinivasan, and D. Vroon. Automatic memory reductions for RTL model verification. In *International Conference on Computer-Aided Design (ICCAD)*, pages 786–793. ACM, 2006.
123. P. Manolios and D. Vroon. Efficient circuit to CNF conversion. In *Theory and Applications of Satisfiability Testing (SAT)*, volume 4501 of *Lecture Notes in Computer Science*, pages 4–9. Springer, 2007.
124. J. McCarthy and J. Painter. Correctness of a compiler for arithmetic expressions. In *Proceedings Symposium in Applied Mathematics, Volume 19, Mathematical Aspects of Computer Science*, pages 33–41. American Mathematical Society, 1967.
125. K. McMillan. *Symbolic Model Checking*. Kluwer Academic, 1993.
126. O. Meir and O. Strichman. Yet another decision procedure for equality logic. In *Proceedings of the 17-th International Conference on Computer Aided Verification (CAV)*, volume 3576 of *Lecture Notes in Computer Science*, pages 307–320. Springer, 2005.
127. M. Mezard, G. Parisi, and R. Zecchina. Analytic and algorithmic solution of random satisfiability problems. *Science*, 297(5582):812–815, 2002.

128. D. G. Mitchell, B. Selman, and H. J. Levesque. Hard and easy distributions for SAT problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 459–465. AAAI Press, 1992.
129. A. Møller and M. I. Schwartzbach. The pointer assertion logic engine. In *PLDI*, pages 221–231, 2001.
130. J. Møller, J. Lichtenberg, H. R. Andersen, and H. Hulgaard. Difference decision diagrams. In *Proceedings of the 13-th International Conference on Computer Science Logic*, volume 1683 of *Lecture Notes in Computer Science*, pages 111–125, 1999.
131. R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky. Determining computational complexity from characteristic phase transitions. *Nature*, 400(8):133–137, 1999.
132. U. Montanari. Networks of constraints: fundamental properties and applications to picture processing. *Information Science*, 7, 1976.
133. M. W. Moskewicz, C. F. Madigan, Y. Zhao, L. Zhang, and S. Malik. Chaff: Engineering an efficient SAT solver. In *Proceedings of the 38-th Design Automation Conference (DAC)*, June 2001.
134. G. Nelson. Techniques for program verification. Technical report, Xerox Palo Alto Research Center, CSL-81-10, 1981.
135. G. Nelson. Verifying reachability invariants of linked structures. In *POPL*, pages 38–47, 1983.
136. G. Nelson and D. C. Oppen. Fast decision procedures based on UNION and FIND. In *Proceedings of the 18-th Annual Symposium on Foundations of Computer Science*, 1977.
137. G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, Oct 1979.
138. G. Nelson and F. F. Yao. Solving reachability constraints for linear lists. Technical report, 1982.
139. G. Nemhauser and L. Wolsey. *Integer and Combinatorial Optimization*. Wiley-Interscience, 1999.
140. R. Nieuwenhuis and A. Oliveras. DPLL(T) with exhaustive theory propagation and its application to difference logic. In *17-th International Conference on Computer Aided Verification (CAV)*, volume 3576 of *Lecture Notes in Computer Science*, pages 321–334. Springer, 2005.
141. R. Nieuwenhuis and A. Oliveras. Proof-producing congruence closure. In *Proceedings of the 16-th International Conference on Term Rewriting and Applications, RTA'05*, volume 3467 of *Lecture Notes in Computer Science*, pages 453–468, Nara, Japan, 2005. Springer.
142. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL( $t$ ). *J. ACM*, 53(6):937–977, 2006.
143. E. Nudelman, K. Leyton-Brown, H. H. Hoos, A. Devkar, and Y. Shoham. Understanding random SAT: Beyond the clauses-to-variables ratio. In *Principles and Practice of Constraint Programming (CP)*, pages 438–452, 2004.
144. Y. Oh, M. N. Mneimneh, Z. S. Andraus, K. A. Sakallah, and I. L. Markov. AMUSE: A minimally-unsatisfiable subformula extractor. In *DAC '04*, pages 518–523, 2004.
145. D. C. Oppen. Complexity, convexity and combinations of theories. *Theoretical Computer Science*, 12(3):291–302, 1980.

146. M. J. Parkinson and G. M. Bierman. Separation logic and abstraction. In *Proceedings of the 32-nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 247–258. ACM, 2005.
147. G. Parthasarathy, M. K. Iyer, K.-T. Cheng, and L.-C. Wang. An efficient finite-domain constraint solver for circuits. In *Design Automation Conference (DAC)*, pages 212–217. ACM Press, 2004.
148. P. F. Patel-Schneider. DLP system description. In *Proceedings DL-98*, pages 87–89, 1998.
149. D. A. Patterson and J. L. Hennessy. *Computer Architecture: A Quantitative Approach*. Morgan Kaufmann, 1990.
150. K. Pipatsrisawat and A. Darwiche. A lightweight component caching scheme for satisfiability solvers. In *Proceedings of 10-th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 294–299, 2007.
151. K. Pipatsrisawat and A. Darwiche. RSAT 2.0: SAT solver description, 2007. SAT solvers competition.
152. D. A. Plaisted and S. Greenbaum. A structure-preserving clause form translation. *Journal of Symbolic Computation*, 2(3):293–304, September 1986.
153. A. Pnueli, Y. Rodeh, and O. Shtrichman. Range allocation for equivalence logic. In *Proceedings of the 21-st Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, Dec 2001.
154. A. Pnueli, Y. Rodeh, O. Shtrichman, and M. Siegel. Deciding equality formulas by small-domains instantiations. In *Proceedings of the 11-th International Conference on Computer Aided Verification (CAV)*, volume 1633 of *Lecture Notes in Computer Science*. Springer, 1999.
155. A. Pnueli, Y. Rodeh, O. Strichman, and M. Siegel. The small model property: How small can it be? *Information and Computation*, 178(1):279–293, Oct. 2002.
156. A. Pnueli and O. Strichman. Reduced functional consistency of uninterpreted functions. In *Pragmatics of Decision Procedures for Automated Reasoning (PDPAR)*, number 898 in *Electronic Notes in Computer Science*, 2005.
157. A. Podelski and T. Wies. Boolean heaps. In *Static Analysis, 12-th International Symposium (SAS)*, volume 3672 of *Lecture Notes in Computer Science*, pages 268–283. Springer, 2005.
158. V. Pratt. Two easy theories whose combination is hard. Technical report, Massachusetts Institute of Technology, 1977.
159. M. Presburger. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In *Comptes Rendus du I congrès de Mathématiciens des Pays Slaves*, 1929.
160. W. Pugh. The Omega test: a fast and practical integer programming algorithm for dependence analysis. *Communications of the ACM*, pages 102–114, 1992.
161. M. O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, July 1969.
162. S. Ranise and C. Tinelli. The SMT-LIB Standard: Version 1.2. Technical report, Department of Computer Science, The University of Iowa, 2006. Available at [www.SMT-LIB.org](http://www.SMT-LIB.org).
163. T. W. Reps, S. Sagiv, and R. Wilhelm. Static program analysis via 3-valued logic. In *Computer Aided Verification, 16-th International Conference, (CAV)*, volume 3114 of *Lecture Notes in Computer Science*, pages 15–30. Springer, 2004.



164. J. C. Reynolds. Reasoning about arrays. *Communications of the ACM*, 22(5):290–299, 1979.
165. J. C. Reynolds. Separation logic: A logic for shared mutable data structures. In *17-th IEEE Symposium on Logic in Computer Science (LICS)*, pages 55–74. IEEE Computer Society, 2002.
166. J. A. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, January 1965.
167. Y. Rodeh. *Techniques in Decision Procedures for Equality Logic and Improved Model Checking Methods*. PhD thesis, Weizmann Institute of Science, 2003.
168. Y. Rodeh and O. Shtrichman. Finite instantiations in equivalence logic with uninterpreted functions. In *Computer Aided Verification (CAV)*, 2001.
169. M. Rozanov and O. Strichman. Generating minimum transitivity constraints in P-time for deciding equality logic. In *Satisfiability Modulo Theories (SMT)*, 2007.
170. H. Ruess and N. Shankar. Solving linear arithmetic constraints. Technical Report SRI-CSL-04-01, SRI, 2004.
171. L. Ryan. Efficient algorithms for clause-learning SAT solvers. Master’s thesis, Simon Fraser University, 2004.
172. A. Sabharwal, C. Ansótegui, C. P. Gomes, J. W. Hart, and B. Selman. QBF modeling: Exploiting player symmetry for simplicity and efficiency. In *Proceedings of 9-th International Conference on Theory and Applications of Satisfiability Testing (SAT)*, pages 382–395, 2006.
173. S. Sagiv, T. W. Reps, and R. Wilhelm. Parametric shape analysis via 3-valued logic. In *POPL*, pages 105–118, 1999.
174. A. Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, 1998.
175. R. Sebastiani. Lazy satisfiability modulo theories. *Journal on Satisfiability, Boolean Modeling and Computation, JSAT*, 3, 2007.
176. B. Selman, H. Levesque, and D. Mitchell. GSAT – a new method for solving hard satisfiability problems. In *AAAI-92*, 1992.
177. S. Seshia, S. Lahiri, and R. Bryant. A hybrid SAT-based decision procedure for separation logic with uninterpreted functions. In *Proc of Design Automation Conference (DAC’03)*, pages 425–430, 2003.
178. R. Shostak. An algorithm for reasoning about equality. *Communications of the ACM*, 21(7):583–585, July 1978.
179. R. E. Shostak. Deciding combinations of theories. *Journal of the ACM*, 31(1):1–12, January 1984.
180. O. Shtrichman. Sharing information between instances of a propositional satisfiability (SAT) problem, Dec 2000. US Patent 2002/0123867 A1.
181. O. Shtrichman. Pruning techniques for the SAT-based bounded model checking problem. In *Proceedings of the 11-th Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARME)*, Sept. 2001.
182. J. Silva and K. Sakallah. GRASP – a new search algorithm for satisfiability. Technical Report TR-CSE-292996, University of Michigan, 1996.
183. E. Singerman. Challenges in making decision procedures applicable to industry. In PDPAR 2005. *Electronic Notes in Computer Science*, 144(2), 2006.
184. L. Stockmeyer. The polynomial-time hierarchy. *Theoretical Computer Science*, 3:1–22, 1977.

185. L. Stockmeyer and A. Meyer. Word problems requiring exponential time. In *Proceedings of the 5-th Annual ACM Symposium on Theory of Computing*, pages 1–9, 1973.
186. O. Strichman. On solving Presburger and linear arithmetic with SAT. In *Formal Methods in Computer-Aided Design (FMCAD)*, pages 160–170, Nov 2002.
187. O. Strichman, S. Seshia, and R. Bryant. Deciding separation formulas with SAT. In *Proceedings of the 14-th International Conference on Computer Aided Verification (CAV)*, volume 2404 of *Lecture Notes in Computer Science*, pages 209–222. Springer, July 2002.
188. A. Stump, C. Barrett, and D. Dill. CVC: a cooperating validity checker. In *Proceedings of the 14-th International Conference on Computer Aided Verification (CAV)*, 2002.
189. A. Stump, C. W. Barrett, D. L. Dill, and J. R. Levitt. A decision procedure for an extensional theory of arrays. In *LICS*, pages 29–37, 2001.
190. A. Stump and L.-Y. Tan. The algebra of equality proofs. In *RTA*, pages 469–483, 2005.
191. M. Talupur, N. Sinha, O. Strichman, and A. Pnueli. Range-allocation for separation logic. In *Proceedings of the 16-th International Conference on Computer Aided Verification (CAV)*, volume 3114 of *Lecture Notes in Computer Science*, pages 148–161. Springer, July 2004.
192. C. Tinelli. A DPLL-based calculus for ground satisfiability modulo theories. In *Proc. 8-th European Conference on Logics in Artificial Intelligence*, volume 2424 of *Lecture Notes in Artificial Intelligence*, pages 308–319. Springer, 2002.
193. C. Tinelli and M. T. Harandi. A new correctness proof of the Nelson–Oppen combination procedure. In *Frontiers of Combining Systems: Proceedings of the 1-st International Workshop (Munich, Germany)*, pages 103–120. Kluwer Academic, 1996.
194. C. Tinelli and C. Zarba. Combining nonstably infinite theories. *Journal of Automated Reasoning*, 34(3):209–238, 2005.
195. G. Tseitin. On the complexity of proofs in propositional logics. In *Automation of Reasoning: Classical Papers in Computational Logic 1967–1970*, volume 2. Springer, 1983. Originally published 1970.
196. R. J. Vanderbei. *Linear Programming: Foundations and Extensions*. Kluwer, 1996.
197. J. Whitemore, J. Kim, and K. Sakallah. SATIRE: A new incremental satisfiability engine. In *IEEE/ACM Design Automation Conference (DAC)*, 2001.
198. R. Wilhelm, S. Sagiv, and T. W. Reps. Shape analysis. In *Compiler Construction CC*, volume 1781 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2000.
199. H. P. Williams. Fourier–Motzkin elimination extended to integer programming problems. *Journal of Combinatorial Theory*, (21):118–123, 1976.
200. R. Williams, C. P. Gomes, and B. Selman. Backdoors to typical case complexity. In *IJCAI*, pages 1173–1178. Morgan Kaufmann, 2003.
201. J. Wilson. Compact normal forms in propositional logic. *Computers and Operations Research*, 90:309–314, 1990.
202. S. A. Wolfman and D. S. Weld. The LPSAT engine & its application to resource planning. In *IJCAI*, pages 310–317, 1999.



203. P. Wolper and B. Boigelot. On the construction of automata from linear arithmetic constraints. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1785 of *Lecture Notes in Computer Science*, pages 1–19. Springer, 2000.
204. L. Wolsey. *Integer Programming*. Wiley-Interscience, 1998.
205. L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown. SATzilla-07: The design and analysis of an algorithm portfolio for SAT. In *Principles and Practice of Constraint Programming (CP)*, pages 712–727, 2007.
206. H. Zantema and J. F. Groote. Transforming equality logic to propositional logic. *Electronic Notes in Computer Science*, 86(1), 2003.
207. L. Zhang, C. Madigan, M. Moskewicz, and S. Malik. Efficient conflict driven learning in a Boolean satisfiability solver. In *ICCAD*, 2001.
208. L. Zhang and S. Malik. Conflict driven learning in a quantified Boolean satisfiability solver. In *Proceedings of International Conference on Computer Aided Design (ICCAD)*, 2002.
209. L. Zhang and S. Malik. Towards symmetric treatment of conflicts and satisfaction in quantified Boolean satisfiability solver. In *Proceedings of 8-th International Conference on Principles and Practice of Constraint Programming (CP)*, 2002.
210. L. Zhang and S. Malik. Extracting small unsatisfiable cores from unsatisfiable Boolean formula. In *Theory and Applications of Satisfiability Testing (SAT)*, 2003.
211. L. Zhang and S. Malik. Validating SAT solvers using an independent resolution-based checker: Practical implementations and other applications. In *Proceedings of the Design and Test in Europe Conference (DATE)*, pages 10880–10885, 2003.

---

# Index

- abstract decision procedure, 237
- abstraction–refinement loop, 73
- Ackermann’s reduction, 66, 162
- adder, 158
  - full adder, 158
- adequate domain, 92
- algorithm, 6
  - ABSTRACTION-REFINEMENT, 74
  - ACKERMANN’S-REDUCTION, 67
  - ADDCLAUSES, 246–249
  - ANALYZE-CONFLICT, 29–38, 50, 247, 249, 251
  - ANTECEDENT, 36, 38
  - ARRAY-REDUCTION, 177
  - BACKTRACK, 30, 31, 248
  - BCP, 30, 31, 247–249
  - BRANCH AND BOUND, 112, 120, 122, 123
  - BRYANT’S-REDUCTION, 70
  - BV-CONSTRAINT, 157, 160, 162
  - BV-FLATTENING, 157
  - CONGRUENCE-CLOSURE, 82, 104
  - DECIDE, 30, 31, 247–249
  - DECIDE-A-CONJUNCTION-OF-EQUALITIES-WITH-EQUIVALENCE-CLASSES, 103
  - DEDUCTION, 245–253, 255, 257, 258, 263, 268
  - DOMAIN-ALLOCATION-FOR-EQUALITIES, 97
  - DPLL( $T$ ), 247–249, 282, 283
  - DPLL( $T$ ), 120
  - DPLL-SAT, 30, 246
  - EAGER-ENCODING, 257, 258
  - EQUALITY-LOGIC-TO-PROPOSITIONAL-LOGIC, 91, 93
  - FEASIBILITY-BRANCH-AND-BOUND, 121, 122
  - GENERAL-SIMPLEX, 117, 121
  - INCREMENTAL BV-FLATTENING, 162
  - LAST-ASSIGNED-LITERAL, 36
  - LAZY-BASIC, 245
  - LAZY-DPLL, 247, 263, 282, 283
  - NELSON-OPPEN, 232
  - NELSON-OPPEN-FOR-CONVEX-THEORIES, 229
  - OMEGA-TEST, 133
  - PRENEX, 212
  - PROJECT, 213, 214
  - QUANTIFIER ELIMINATION, 214
  - REMOVE-CONSTANTS, 60
  - REMOVE-CONSTANTS-OPTIMIZED, 104
  - RESOLVE, 36, 38
  - SAT-SOLVER, 162, 245, 257, 258
  - SEARCH-BASED DECISION OF QBF, 220
  - SEARCH-INTEGRAL-SOLUTION, 121
  - SIMPLIFY-EQUALITY-FORMULA, 86
  - STOP-CRITERION-MET, 36
  - VARIABLE-OF-LITERAL, 36
- algorithm portfolio, 55
- aliasing, 184
- all-SAT, 57
- antecedent, 3
- arithmetic right shift, 156

- arrangement, 237
- array, 171
  - bounds violation, 173
  - extensionality, 175
- array logic, 171, 172
- array property, 175
- assignment, 5
  - full, 5
  - partial, 5
- atom, 8, 83
- automated reasoning, 25
- axioms, 3, 62
  
- back-door variables, 27
- basic variables, 116
- BCP, 31, 42, 251
- BDD, 43, 45, 47, 48, 50, 56, 57, 92, 101, 107, 108, 169, 223, 253, 267
  - APPLY procedure, 46–49, 53
  - multiterminal, 53
  - reduced ordered, 44
  - reduced ordered (ROBDD), 43
- Bellman–Ford algorithm, 142
- binary decision diagrams, *see* BDD
- binary decision tree, 43
- binary encoding, 153
- binary tree, 191
- binding scope, 208
- bit vector, 49, 152
- bit-blasting, 156
- bit-vector, 65
- bit-vector arithmetic, 151, 163
- blocking clause, 242, 246
- Boolean constraints propagation, *see* BCP
- Boolean Fourier–Motzkin, 267
- bounded model checking, 20, 169
- Bryant’s reduction, 66
  
- canonicity, 43
- carry bits, 158
- case-splitting, 11, 107
  - semantic, 11
  - syntactic, 11
- certificate, 223
- chord, 90
- chord-free cycle, 90
- chordal graph, 90
- clause, 12
  - antecedent, 29
  - asserting, 33, 34
  - conflicting, 29, 32, 38
  - satisfied, 29
  - unary, 29, 33
  - unit, 29
    - unresolved, 29
- CNF, 13, 27, 37, 43, 57, 150, 160, 214, 220, 223, 278, 280, 282
- colorability, 26
- compiler, 111
- complete proofs, 255
- completeness, 4, 6, 7, 63
- concatenation, 153
- conflict clause, 33–35, 38, 41, 52, 246, 251
- conflict graph, 32, 34, 38
- conflict node, 32, 34, 36
- conflict-driven backtracking, 33–35, 52
- congruence closure, 82, 107, 241
  - abstract congruence closure, 107
- conjunctive fragment, 16
- conjunctive normal form, *see* CNF
- consequent, 3
- constraint satisfaction problem (CSP), 17
- contradiction, 5
- contradictory cycle, 85, 86
  - simple, 85
- convex theory, 226
- cut
  - separating cut, 34
  - Gomory, *see* Gomory cut
- cutting-planes, 122, 145
  
- Davis–Putnam–Loveland–Logemann, *see* DPLL
- De Morgan’s rules, 8
- decidability, 7, 22, 62
- decision level, 29, 31
- decision problem, 6
- decision procedure, 7
- delayed theory combination, 266
- derivation tree, 12
- difference logic, 140
- DIMACS format, 278
- disequality literals set, 83
- disequality path, 84, 85
  - simple, 85

- disjunctive linear arithmetic, 241
  - disjunctive normal form, *see* DNF
  - DNF, 10, 214
  - domain, 207
  - domain allocation, 92, 96, 101, 106
  - don't care variables, 46
  - DPLL, 28, 29, 51, 54, 246
  - DPLL( $T$ ), *see* algorithms
  - dynamic data structure, 182
  
  - eager encoding, 254
  - ellipsoid method, 113
  - empirical hardness models, 55
  - encoder, 241
  - endianness, 189
  - equality graph, 84–87, 95, 96, 99
    - nonpolar, 89, 91
  - equality literals set, 83
  - equality logic, 16, 59
  - equality path, 84, 97, 99, 104, 250
    - simple, 85
  - equality sign, 226
  - equisatisfiable, 8, 60, 78, 86, 91, 255
  - EUF, 61
  - exhaustive theory propagation, 250
  - existential nodes, 218
  - existential quantification, 176
  - existential quantifier ( $\exists$ ), 207
  - explanation, 251
  - expressiveness, 18
  
  - first UIP, 36
  - first-order logic, 14
  - first-order theory, 2
  - fixed-point arithmetic, 165
    - saturation, 166
  - floating-point arithmetic, 165
  - forall reduction, 215
  - Fourier–Motzkin, 113, 126, 129, 130,
    - 132, 133, 145, 217, 267
  - fragments, 16
  - free variable, 16, 208
  - functional congruence, 62
  - functional consistency, 62, 64, 162
  
  - Gaussian variable elimination, 118
  - general form, 114
  - general simplex, 112, 113
  - Gomory cut, 123
  
  - graph, 272
  - ground formulas, 17
  - ground level, 29, 33
  
  - Hoare logic, 20
  - hyperresolution graph, 41
  
  - ILP, 112, 143, 144, 169
    - 0–1 linear systems, 140
    - relaxation, 120
  - implication graph, 31, 50
  - incomplete procedure, 62
  - incremental, 120
  - incremental satisfiability, 55, 246
  - inequality graph, 142
  - inference rule
    - $>$ -CONTR, 258
    - $>$ -TRANS, 258
    - BINARY RESOLUTION, 37, 252
    - CONSTANTS, 261, 262, 264
    - CONTRADICTION, 3, 4
    - DOUBLE-NEGATION, 4
    - DOUBLE-NEGATION-AX, 3, 4
    - EQ-CONTRADICTION, 256
  - instantiation, 4
  - M.P., 3, 4
  - ORDERING I, 256
  - ORDERING II, 256
  - PROJECT, 261, 262, 264
  - SUBSTITUTION, 256
- inference rules, 3, 254
  - inference system, 3
  - initialized diameter, 221
  - integer linear programming, *see* ILP
  - interpretation, 15
  
  - job-shop scheduling, 141
  
  - languages, 18
  - learning, 11, 33, 42, 51, 218, 219
  - least significant bit, 153
  - lemma, 243
  - lifetime, 183
  - linear arithmetic, 111
  - linear programming, 112
  - linked list, 190
  - literal, 8, 83
    - satisfied, 9
  - logical axioms, 16, 226

- logical gates, 12
- logical right shift, 156
- logical symbols, 226
- loop invariant, 199
  
- matrix, *see* quantification suffix
- maximally diverse, 234
- memory
  - layout, 182
  - location, 182
  - model, 181
  - valuation, 182
- mixed integer programming, 122
- model checking, 56, 73, 223, 224
- model-theoretic, 3
- modular arithmetic, 150
- modulo, 154
- monadic second-order logic, 204
- most significant bit, 153
- multiterminal BDD, *see* BDD
  
- negation normal form, *see* NNF
- Nelson–Oppen, 225, 227, 228, 236, 266
- NNF, 8, 14, 22, 83, 84, 87, 105, 176, 178, 242, 244, 263
- nonchronological backtracking, 218, 219
- nonlinear real arithmetic, 224
- nonlogical axioms, 226
- nonlogical symbols, 226
- normal form, 8
- NP-complete, 209
  
- offset, 189
- Omega test, 113, 129, 267
  - dark shadow, 136
  - gray shadow, 137
  - real shadow, 133
- overflow, 150
  
- parse tree, 271
- parsing, 271, 275
- partial implication graph, 32
- partially interpreted functions, 73
- Peano arithmetic, 17, 223
- phase, 9, 214
- phase transition, 56
- pigeonhole problem, 38
- pivot operation, 118
- pivoting, 118
  
- planning, 210
- planning problem, 25
- plunging, 250
- pointer, 181
- pointer logic, 186
- points-to set, 185
- polarity, 9
- predicate abstraction, 205
- predicate logic, *see* first-order logic
- prenex normal form, 212, 213, 218
- Presburger arithmetic, 16, 171, 211, 223
- procedure, 6
- program analysis, 184
- projection, 213
- proof, 254
- proof constraint, 254
- proof graph, 254
- proof rule, *see* inference rule
- proof step, 254
  - constraint, 254
- proof-theoretic, 3
- propositional encoder, 156
- propositional encoding, 257
- propositional skeleton, 83, 89, 156, 242, 281
- PSPACE-complete, 209
- pure literal, 9
- pure variables, 196
- purification, 228
  
- Q-resolution, 223
- QBF, *see* quantified Boolean formula
- QBF search tree, 218
- quantification prefix, 212
- quantification suffix, 212, 214, 218
- quantified Boolean formula, 209
- quantified disjunctive linear arithmetic, 211
- quantified propositional logic, 209
- quantifier, 207
- quantifier alternation, 176, 207
- quantifier elimination, 213
- quantifier-free fragment, 16
  
- reachability predicate, 198
- reachability predicate formula, 198
- read-over-write axiom, 172
- reduction rules, 44
- reference, 183

- resolution, 54
  - binary resolution, 37, 41, 214, 216
  - binary resolution graph, 41
  - resolution graph, 41, 42
  - resolution variable, 37, 38
  - resolvent clause, 37
  - resolving clauses, 37
- restarts, 42
- restrict operation, 46, 57, 216
- rewriting rules, 74, 106
- rewriting systems, 74
- ripple carry adder, 158
- rounding, 167
- routing expressions, 204
  
- SAT decision heuristic, 39
  - Berkmin, 40
  - conflict-driven, 40
  - DLIS, 40
  - Jeroslow–Wang, 39
  - VSIDS, 40
- SAT solvers, 27, 267
- satisfiability, 5
- saturation, 260
- semantics, 5
- sentence, 16, 208, 226
- separation logic, 146, 192, 205
- Shannon expansion, 47, 223
- shape analysis, 205
- side conditions, 261
- sign bit, 154
- sign extension, 155
- signature, 16
- simplex, 11, 112, 241
  - basic variable, 116
  - nonbasic variable, 116
  - additional variable, 114
  - Bland’s rule, 119
  - pivot column, 118
  - pivot element, 118
  - pivot operation, 117
  - pivoting, 118
  - problem variable, 114
- Skolem variable, 202
- Skolemization, 202
- small-domain instantiation, 92, 267
- small-model property, 4, 92, 122, 205
- SMT-COMP, 241, 265, 269
- SMT-LIB, 269
  
- sort, 59
- soundness, 4, 6, 7, 63, 98
- sparse method, 88
- state space, 92–96, 98, 106
- static analysis, 185
- static-single-assignment, 20, 63
- stochastic search, 28
- structure in formal logic, 15
- structure type, 189
- subsumption, 14
- symbol table, 278
- symbolic access paths, 204
- symmetric modulo, 131
  
- T*-satisfiable, 16, 226
- T*-valid, 16, 226
- tableau, 116
- tautology, 5
- term, 10
- theory, 2, 15, 61
- theory combination, 226
- theory combination problem, 226
- theory of equality, 59
- theory propagation, 243, 248, 255
- timed automata, 146
- tools
  - BARCELOGIC, 265
  - BERKMIN, 41, 55
  - C32SAT, 169
  - CHAFF, 40, 54
  - COGENT, 169
  - CVC, 169, 264, 265
  - CVC-LITE, 265
  - DLSAT, 265
  - GRASP, 54
  - GSAT, 54
  - ICS, 169
  - ICS-SAT, 265
  - LASH, 224
  - MATHSAT, 265
  - MINISAT, 55
  - MINISAT-2, 55
  - SIEGE, 55
  - SLAM, 205
  - SPEAR, 170
  - STP, 169
  - SVC, 169
  - TERMINATOR, 205
  - UCLID, 267

- VERIFUN, 265
- WALKSAT, 54
- YICES, 169, 265
- total order, 19
- transitive closure, 193, 204
- translation validation, 75, 77
- truth table, 4, 5
- Tseitin's encoding, 12–14, 21, 160
- Turing machine, 74
- two's complement, 153, 154, 156
- two-player game, 209
- type, 152
- type checking, 272
- UIP, 36
- unbounded variable, 128
- uninterpreted functions, 60–67, 69–75, 77–83, 104, 106, 107, 225, 228, 229
- uninterpreted predicates, 61
- union-find, 82, 107
- unique implication point, *see* UIP
- unit clause rule, 29
- universal nodes, 218
- universal quantification, 176
- universal quantifier ( $\forall$ ), 207
- unsatisfiable core, 42, 252, 259
- validity, 5
- verification condition, 20, 26, 172, 199
- well-formed, 14
- winning strategy, 209
- write rule, 174
- zero extension, 155
- $\lambda$ -notation, 151
- $\Sigma$ -formula, 16
- $\Sigma$ -theory, 16