

P. Madhusudan, Mahesh Viswanathan

Logic in Computer Science

Rough course notes

January 20, 2026

Contents

| | |
|---|----|
| 0 Logic over Structures: A Single Known Structure, Classes of Structures, and All Structures | 1 |
| 0.1 Logic on a Fixed Known Structure | 1 |
| 0.2 Logic on a Fixed Class of Structures | 3 |
| 0.3 Logic on All Structures | 5 |
| 0.4 Logics over structures: Theories and Questions | 7 |
| 1 Propositional Logic | 11 |
| 1.1 Syntax | 13 |
| 1.2 Semantics | 16 |
| 1.3 Satisfiability and Validity | 19 |
| 1.4 Compactness Theorem | 26 |
| 1.4.1 Compactness using König's Lemma | 26 |
| 1.4.2 Compactness using Henkin Models | 29 |
| 1.4.3 An Application of Compactness: Coloring Infinite Planar Graphs | 30 |
| 2 Proof Systems | 33 |
| 2.1 A Frege-style Proof System | 34 |
| 2.1.1 Completeness Theorem | 36 |
| 2.2 Resolution | 37 |
| 2.2.1 Proving Tautologies with Resolution | 40 |
| 2.2.2 Completeness of Resolution | 42 |
| 2.3 Craig's Interpolation Theorem and Proof Complexity | 44 |
| 2.3.1 Craig's Interpolation Theorem | 44 |
| 2.3.2 Size of Interpolants | 45 |
| 2.3.3 Interpolants from Refutations | 48 |
| 2.3.4 Lower bounds on Resolution Refutations | 54 |

| | | |
|----------|---|-----|
| 3 | First Order Logic | 59 |
| 3.1 | Syntax | 59 |
| 3.2 | Semantics | 61 |
| 3.2.1 | Satisfiability, Validity, and First order theories | 65 |
| 3.3 | Overview | 66 |
| 4 | Quantifier Elimination and Decidability | 69 |
| 4.1 | Dense Linear Orders without Endpoints | 71 |
| 4.2 | Linear Arithmetic | 74 |
| 4.2.1 | Fourier-Motzkin | 75 |
| 4.2.2 | Ferrante-Rackoff | 78 |
| 4.3 | Other theories that admit quantifier elimination | 80 |
| 5 | Lower Bounds for the Validity Problem | 81 |
| 5.1 | Number Lines | 82 |
| 5.2 | Church-Turing Theorem | 84 |
| 5.3 | Trakhtenbrot's Theorem | 88 |
| 6 | Incompleteness Theorems | 91 |
| 6.1 | Gödel's (First) Incompleteness Theorem | 92 |
| 6.2 | Incompleteness of the theory of natural numbers with additional and multiplication | 93 |
| 6.3 | Program Verification | 96 |
| 6.4 | Further Remarks | 97 |
| 7 | Quantifier-free theory of equality | 101 |
| 7.1 | Decidability using Bounded Models | 101 |
| 7.2 | An Algorithm for Conjunctive Formulas | 102 |
| 7.2.1 | Computing $CC(E)$ | 105 |
| 7.3 | Axioms for The Theory of Equality | 107 |
| 8 | Completeness Theorem: FO Validity is r.e. | 111 |
| 8.1 | Prenex Normal Form | 113 |
| 8.2 | Skolemization / Herbrandization | 113 |
| 8.3 | Herbrand's theorem | 115 |
| 8.4 | Some consequences of Herbrand's theorem | 120 |
| 8.5 | Gödel's completeness theorem: FO Validity is recursively enumerable | 121 |
| 8.5.1 | The case of finite sets of formulas | 124 |
| 8.5.2 | The case for infinite sets of formulas | 125 |
| 8.5.3 | Completeness Theorem | 126 |
| 8.6 | Observations and Consequences | 128 |

| | |
|--|-----|
| A Computability and Complexity Theory | 131 |
| A.1 Turing Machines | 132 |
| A.2 Church-Turing Thesis | 135 |
| A.3 Recursive and Recursively Enumerable Languages | 137 |
| A.4 Reductions | 139 |
| A.5 Complexity Classes | 142 |
| A.6 Relationship between Complexity Classes | 145 |
| A.7 P and NP | 147 |
| A.7.1 Alternate characterization of NP | 148 |
| A.7.2 Reductions, Hardness and Completeness | 149 |

Chapter 0

Logic over Structures: A Single Known Structure, Classes of Structures, and All Structures

In this chapter, we give a very informal account of first-order logic over structures (a single structure, a class of structures, and all structures), formulate some questions, and look at some important results we will prove eventually in this book.

0.1 Logic on a Fixed Known Structure

You are probably already familiar with first-order logic on a fixed structure (or context or world). For example, you perhaps know what this statement means:

$$\forall x \in \mathbb{N} \ \exists y \in \mathbb{N} \ x < y$$

It says “for every natural number x , there exists a natural number y such that $x < y$. You have perhaps learned this in elementary courses on discrete mathematics.

If you haven’t or need a primer, I recommend reading it from the following sources:

- Madhu’s primer for CS173: <https://courses.grainger.illinois.edu/cs173/fa2017/B-lecture/NotesByMadhu/Notes-1.pdf>

The main thing to note here is that you *know* the structure/universe/context you are talking about—in this case natural numbers.

You should make sure you know several things about such logical notation:

- You should know the Boolean connectives \wedge , \vee , and \neg .
- You should know the meaning of \forall (“forall”) and \exists (“exists”), which are *quantifiers*.
- You should know that $\alpha \Rightarrow \beta$ (read α *implies* β) has a formal meaning that is precisely the same as $(\neg\alpha) \vee \beta$, and not some other uses in English. For example, “Goldilocks is the president of the United States implies all gorillas are green” is a statement that is true in the current world, since Goldilocks is not the president of the United States (I am assuming this is true when you read this as well).

2 0 Logic over Structures: A Single Known Structure, Classes of Structures, and All Structures

- One really needs only the connectives \vee and \neg ; the rest can be derived or defined as shortcuts:
 - $\alpha \wedge \beta$ is logically equivalent to $\neg(\neg\alpha \vee \neg\beta)$
 - $\alpha \Rightarrow \beta$ is logically equivalent to $(\neg\alpha) \vee \beta$
 - $\alpha \Leftrightarrow \beta$ is logically equivalent to $(\alpha \Rightarrow \beta) \wedge (\beta \Rightarrow \alpha)$.
- You should know the de Morgan laws and how negation goes into quantifiers:
 - $\neg(\alpha \vee \beta)$ is equivalent to $(\neg\alpha) \wedge (\neg\beta)$
 - $\neg(\alpha \wedge \beta)$ is equivalent to $(\neg\alpha) \vee (\neg\beta)$
 - $\neg(\exists x. \alpha)$ is equivalent to $\forall x. (\neg\alpha)$
 - $\neg(\forall x. \alpha)$ is equivalent to $\exists x. (\neg\alpha)$

Using the above, you should know that you can “push” the negations all the way in so that they are applied only to atomic symbols/formulae.

As a concrete example, let us discuss logics over a fixed structure—the set of natural numbers \mathcal{N} .

In propositional logic over a fixed structure (or universe or world), you would have a mapping between propositional symbols and *statements*. For example, p could be the statement “there are finitely many primes” and q could be the statement “all primes other than 2 are odd”. Then p is false in this world, and q is true in this world, etc.

First order logic over natural numbers is more powerful and useful. Over natural numbers we have several functions and relations that we know. For example, $+$ is a (binary) function, \times is a (binary) function, and *square* is a (unary) function. And $<$, \leq are all (binary) relations. In fact, another relation that we overlook sometimes is the *equality* relation $=$, since it’s so common.

Now, let us fix this signature of *symbols* for functions and relations. Let us also fix a set of symbols to denote *variables*, called *Var*. And define a first order logic formally:

Terms: $t, t' ::= x \mid c \mid +(x, y) \mid \times(x, y) \mid \text{square}(x)$

Formulas: $\varphi, \varphi' ::= t = t' \mid t < t' \mid t \leq t' \mid \varphi \vee \varphi' \mid \varphi \wedge \varphi' \mid \neg\varphi \mid \forall x. \varphi \mid \exists x. \varphi$

where $c \in \mathbb{N}, x \in \text{Var}$.

We have here two kinds of expressions—terms and formulas. Terms are obtained from constants and variables by applying functions, recursively, and intuitively “evaluate” to some number. For example, $+(5, 8)$ is a term. Note that we write functions like $+$ with the symbol in front, rather than in infix notation; i.e., we write $+(5, 8)$ instead of $5+8$. But as you are familiar with computer science, you can think of $+$ as a function that you “call”, and so $+(5, 8)$ should make sense. $+(9, x)$ is a term as well. And $\text{square}(*(\text{square}(13384398), \text{square}(y)))$ is a term too. Clearly there are infinitely many terms.

We don’t really need all constants in the grammar. 0 and 1 are sufficient, as any other constant can be expressed as an appropriate sum of 1s.

Formulas, on the other hand, evaluate to a Boolean, i.e., *true* or *false*. The *atomic* formulas are those that are formed from terms— those of the form $t = t'$, $t < t'$ and $t \leq t'$. Formulas are basically atomic formulas closed under Boolean operations and quantification.

Note that we write $\forall x. \exists y. x < y$, rather than $\forall x \in \mathbb{N}. \exists y \in \mathbb{N}. x < y$, since the universe over which variables are quantified are assumed to be natural numbers anyway. If there are multiple sorts of objects that we want to quantify over, we would introduce the more general syntax where every time we quantify, we will say which sort it is over (or designate different sets of variables for each sort). This is similar to programming, where we declare variables to be of different sorts— like integers or strings.

You should be able to read and write FO formulas over natural numbers. For instance, the formula $\forall x. (x > 2) \Rightarrow \text{square}(x) > x$ says “the square of any number greater than two is larger than it”, which is happens to be true over natural numbers. Similarly, if we wanted to say “every number other than 0 has a smaller number”, we would write this as $\forall x. \neg(x = 0) \Rightarrow (\exists y. (y < x))$, which also happens to be true over natural numbers.

Now consider the formula $\forall x. (x > y)$. Is it true over natural numbers? It’s hard to imagine how to interpret the statement as its unclear what y is. This leads us to define something called sentences.

A *sentence* is a formula where every variable is quantified, We will define this formally later, but it should be clear to you what this means.... any variable that occurs in the formula must have a quantification “outside” it such that the variable is in scope of that quantifier. This is similar to programming— we may require all variables used in a program to be *declared*, i.e., every use of a variable should be in the scope of a declaration of it. Sentences evaluate to true or false on a structure.

The first-order *theory* of the structure \mathbb{N} is the set of FO sentences that hold in that model, denoted $\text{Th}(\mathbb{N})$. Note that since any sentence α must either be true or false in the structure, and hence either α or $\neg\alpha$ must be in the theory. More generally, a theory is just a set of sentences. And a theory is said to be *complete* if for every sentence α , either α or $\neg\alpha$ belongs to it. So $\text{Th}(\mathbb{N})$ is complete (indeed, the theory of any fixed structure is complete).

0.2 Logic on a Fixed Class of Structures

In mathematics and computer science, we often want to express properties that hold on a *class* of structures, not just a single structure. For example, we may ask what formulas/sentences are true over *groups*, or over *finite graphs*, or over *trees*, or over *linked lists*, or over *recursively defined datatypes*, or over *SQL (relational) databases*, or about *objects* in a class.

Unlike a single structure, like \mathbb{N} or \mathbb{Z} or \mathbb{R} , we are interested in a *class* of structures.

Groups

Let us take groups. A group is defined as a set S endowed with a binary relation $\circ : S \times S \rightarrow S$ that satisfies the following properties:

Associativity: for every $a, b, c \in S$, $(a \circ b) \circ c = a \circ (b \circ c)$

Identity: There is an element $e \in S$ such that for every $a \in S$, $a \circ e = e \circ a = a$.

Inverse: For every $a \in S$, there exists an $a' \in S$ such that $a \circ a' = a' \circ a = e$.

For example, the set of integers \mathbb{Z} with the operator $+$ forms a group (0 is the identity, and for every $i \in \mathbb{Z}$, $-i$ is its inverse. Another class of examples of groups is obtained by taking a finite set E and considering the set of elements consisting of permutations of E , with the binary operation being *composition* of permutations (a composition of permutations is a permutation as well). The identity permutation is the identity element, and every permutation has an inverse, of course, which “reverts” the permutation.

Now, there are of course properties that hold on *all* groups. For instance, the identity element must be unique. Here’s a proof: Assume $e, f \in S$ are both identity elements. Then $e.f = e$ (since f is an identity) and $e.f = f$ (since e is an identity). Hence $e = f$.

Similarly, one can show that every element’s inverse is unique.

The above proof shows that the fact that 0 is the unique identity element for $+$ over \mathbb{Z} is not a particular property satisfied only on integers, but rather is a property shared among all groups. The field of group theory studies groups in their own right, since they occur commonly in many areas and applications.

We can now define the *first order theory* of groups as the set of all FO sentences that hold over groups. The signature could include a special constant e to denote the identity element. The sentence $\forall a, b, c. (a \circ b = e \wedge b \circ a = e \wedge a \circ c = e \wedge c \circ a = e) \Rightarrow b = c$, which says that every element has a unique inverse, is hence a theorem in this theory.

Note however that the theory of groups is not a complete theory. For example, the sentence $\forall x, y. x \circ y = y \circ x$ is not a theorem nor is its negation a theorem. There are some groups where this property is true (like $+$ over integers) and some where it’s not true (like permutations of a fixed finite set of elements).

Graphs

Graphs are ubiquitous in computer science. Each graph can be seen as a model where there is a set/universe which is *finite* and that has a *binary relation* E over it, which models the set of edges. We would expect E to be *symmetric* (if $E(u, v)$ holds, then $E(v, u)$ also holds). And we don’t want “self-edges” (for any u , $E(u, u)$ does not hold).

One can then defined the *theory of graphs*— which consists of all FO sentences true over graphs. This FO theory is not terribly interesting, as there are very few properties about graphs you can express using just FO theory on graphs. Note that

this theory is not complete, again, of course—the sentence $\forall u, v. E(u, v)$ is neither true in all graphs nor false in all graphs.

One can of course define any subclass of graphs—such as planar graphs or bipartite graphs—and talk about the theory of such subclasses as well. How does the theory of graphs, TG , and the theory of planar graphs, TPG , compare? Is one a subset of the other? It is easy to see that any sentence that holds for all graphs is certainly true for all planar graphs as well. Hence $TG \subseteq TPG$.

In general, if C and \mathcal{D} are two classes of structures with C a subclass of \mathcal{D} , then the theory of \mathcal{D} would be a subset of the theory of C . The smaller the class, the larger its theory! In the limit, when there is only one structure, the theory becomes complete (and of course cannot get any larger without containing contradictions).

One can now ask—do we really need to have a single structure in order to obtain a complete theory? In other words, is there a class of structures/models C that has at least two structures such that its first order theory is complete? Strangely, the answer is *yes*! We will see examples of this in the course. For example:

- The theory of rationals with only the relations $=$ and \leq is the same as the theory of reals with the relations $=$ and \leq ! In other words, there is no *first-order sentence* that can distinguish between these structures.

Note that the above is very specific to the fact that we have only FOL and only the fixed signature involving \leq . For example, if we had the function symbol *square* that returns the square of a rational/real, then we *can* distinguish the two structures. (Can you come up with one? If you have in addition constants such as $0, 1, 2, \dots$, it would be simpler.)

- More generally, the theory of dense linear orders without endpoints (no least or largest element) is complete. No matter which dense linear order you pick, you will find that the theory is the same! So here is an example of an infinite class of structures which no first order formula (with only \leq in the signature) can distinguish.
- One extremely surprising result is that there are structures that are *non-isomorphic* to natural numbers and yet satisfy the same first-order properties of natural numbers! You can imagine an alien species having such a *non-standard* model of arithmetic in their head (though I wonder what kind of evolutionary circumstance would give rise to such models in their psyche), and yet we would agree with them about all theorems in FOL over arithmetic!

0.3 Logic on All Structures

Finally, we can consider logics on *all structures*. This may sound a bit unnatural and not very useful. But as we shall see, it is quite useful, as it gives a way to study general metatheorems in logic that are independent of a particular structure or class of structures.

One reason why logics over all possible structures is that one can *carve* out useful and natural fragments using *axiomatizations*. Axiomatizations are *logical* mechanisms of specifying a class of structures that you want to study.

Axiomatizations are, in a certain sense, the purest form of reasoning about a class of structures. If we want to reason about a class of structures $\{C\}$, then it remains how to *define* them formally so that we can reason with them. For instance, you and I may think we know what natural numbers and arithmetic are, but to formally reason with arithmetic, we must be able to state our assumptions clearly. If you start an argument with “There are infinitely many even numbers, and ...” and I interrupt you and ask you why that is so and I don’t believe it, then you may provide a proof of it. But every proof you give will make assumptions (hopefully simpler assumptions), till at some point you give up and say that those are self obvious. For example, you may refuse to give a proof of why $x + 0 = x$. So it is natural to ask whether there are some fundamental and simple assumptions about natural numbers that we all can agree upon (and people who don’t believe them can go climb a gum tree) such that all arguments can be made only using such assumptions.

Presburger arithmetic is a particular set of axioms that characterize natural numbers with addition only. The axioms are:

- (A1) $\forall x. \neg(x + 1 = 0)$
- (A2) $\forall x, y. (x + 1 = y + 1) \Rightarrow x = y$
- (A3) $\forall x. x + 0 = x$
- (A4) $\forall x, y. x + (y + 1) = (x + y) + 1$
- (A5) For any first order formula P with a free variable x , the following holds:

$$(P(0) \wedge \forall x. (P(x) \Rightarrow P(x + 1))) \Rightarrow \forall y. P(y)$$

The meaning and soundness of axioms (A1)–(A4) should be obvious. (A5) is actually a *set* of axioms, called an axiom schema. It says that for any property P of natural numbers expressible in FOL, induction is a sound way of proving that P holds on all natural numbers. More precisely, if it was true that P holds for 0 and for every x , if P holds for x then P holds for $x + 1$, then P must hold for all natural numbers.

It is an amazing fact that all first-order properties of natural numbers with addition can be proved just using the axioms above. We will not show this in this course, however, but argue a related theorem to show that the theory of natural numbers with addition is decidable. (It will become clear later why these are related.)

Perhaps a more natural example of axiomatizations is the characterization of groups. Recall that groups satisfy a specific set of properties (associativity, existence of identity, and existence of inverses). This is *the* definition of groups—we do not have some other “natural” model of groups in our minds. And furthermore it turns out that we can characterize the properties of groups in FOL, with the signature containing $=$ and the sole binary function \circ , and an identity element (constant symbol) e :

- (A1) Associative law: $\forall x, y, z. x \circ (y \circ z) = (x \circ y) \circ z$

(A2) Identity: $\forall x. (x \circ e = x \wedge e \circ x = x)$

(A3) Existence of inverses: $\forall x. \exists y. (x \circ y = e \wedge y \circ x = e)$

Note that in the above, the set of axioms is *finite* as opposed to Presburger arithmetic. Similar to the above, we can characterize many classes of algebraic structures—abelian groups, rings, fields, Boolean algebras, etc.

Given a set of axioms \mathcal{A} , we can think of it as culling out a class of structures from the class of all structures, namely those that satisfy the axioms. We can then talk about the *theory* of the axioms—the set of sentences that are satisfied in every structure in this culled out class of structures. Let $Th(\mathcal{A})$ denote the theory defined by the structures that satisfy the axioms \mathcal{A} .

For example, for the set of axioms defining groups above, its theory contains the statement: $\forall x, y, z. (x \circ y = e \wedge y \circ x = e \wedge x \circ z = e \wedge z \circ x = e) \Rightarrow (y = z)$

0.4 Logics over structures: Theories and Questions

Given the above discussion, we have four kinds of theories:

- The theory of a single structure M , denoted $Th(M)$.
- The theory of a class of structures C , denoted $Th(C)$.
- The theory of all structures, which we will call tautologies or valid sentences, denoted $Th(FOL)$.
- The theory of a set of axioms $Th(\mathcal{A})$, which is the theory of the class of structures that satisfy \mathcal{A} .

Given the above, one can ask many interesting questions and make some simple observations:

- Let us call a theory complete if for every sentence α , either α is in the theory or $\neg\alpha$ is in the theory.
- Note that the theory of a single structure M is always complete. However, the theory of a class of structures need not be complete. (For example, over groups, the sentence $\forall x, y. x = y$ is neither in the theory nor is its negation in the theory). Similarly, the theory of a set of axioms need not be complete.
- The theory of a structure or a nonempty class of structures can never have a contradiction—i.e., it cannot have both α and $\neg\alpha$. You can have one only or the other or neither, but not both. The theory of the empty class of structures contains all sentences, and hence contradictions.
- The theory of a set of axioms can have a contradiction. But this happens only if they define an *empty* class of structures. Otherwise, the theory will have no contradiction. (For example, if axioms include $\forall x. P(x)$ and $\exists x. \neg P(x)$, then the set of structures defined by these axioms is empty, and its theory contains all sentences.)

8 0 Logic over Structures: A Single Known Structure, Classes of Structures, and All Structures

- If a (nonempty) class of structures C is a subclass of a class of structures \mathcal{D} , then $Th(\mathcal{D}) \subseteq Th(C)$. So as the class of structures get smaller, its theory gets larger. (In the limit, the class of structures is a single structure and the theory becomes largest and is complete.)
- What is the complexity of deciding whether a sentence is valid (on all structures)? Is it decidable? If not, is it recursively enumerable? What is its precise complexity?
- Is there a set of axioms (some regular simple set of axioms) that characterize natural numbers with addition? Integers with addition? Natural numbers or integers with addition and multiplication? What about rationals with addition and/or multiplication? Reals?
- Is it possible to decide if a set of axioms has a contradiction? I.e., whether there is at least one structure satisfying it?
- For each of the above theories, independent of whether they can be axiomatized or not, are the theories decidable? (Can we build programs that check whether a theorem is true or not, i.e., belongs to the theory or not, completely automatically?)
- How do we know that proofs even exist? Can it be the case that there are (natural) classes of structures for which some theorems do not have proofs?
- If we fix a set of axioms, is it true that every theorem (statement in its theory) has a proof, always, that follow from the axioms? Maybe first we need to define what a proof is? What's a proof?

I encourage you to think of all such combinations of questions. Some will be trivial and you will be able to answer them. Most others you will be able to answer at the end of the course. These questions that will occupy us for roughly half of this course.

Here is a sample of the remarkable theorems in logic you will learn in this course:

- The set of all valid sentences (the set of sentences that are true in all structures) is not decidable. However, it is recursively enumerable!
- One can in fact set up a formal proof mechanism for proving valid sentences. Proofs are syntactic objects that (a) are finite and (b) can be verified easily using syntactic rules. And then we can show that any valid sentence has a proof! (This is Gödel's completeness theorem.) It then follows that a Turing machine/program can simply look for such proofs, and hence the set of valid formulas is recursively enumerable.
- More remarkably, if I have a set of axioms \mathcal{A} where \mathcal{A} is a finite set of axioms, then the theory of \mathcal{A} is also recursively enumerable. This even holds if \mathcal{A} is infinite and is a computable (or even recursively enumerable set). In fact, one can set up generic proof systems that work by assuming any set of axioms that can prove any theorem in the theory of the set of axioms. This is remarkable as it shows that any axiomatizable theory has proofs and a computer can just look for such proofs!
- If a set of axioms define a *complete* theory, then the theory is even *decidable*! There is a program that can take a statement and decide whether it is a theorem or not. (Don't ask me how long this will take, though!)

- There are several specific structures and signatures where you can ask whether its theory is decidable (naturals, integers, rationals, reals with addition and/or multiplication, etc.). Most of these have been settled. Note that a complete axiomatization of them also implies decidability, and hence undecidability of the theory means there is no complete axiomatization.

Four remarkable results are:

- The theory of natural numbers with addition is decidable (Presburger arithmetic above is a complete axiomatization).
- The theory of natural numbers with addition and multiplication is *not* decidable. It follows that this theory hence does not have a complete axiomatization. This is essentially one of Gödel's incompleteness theorems. It turns out that even validity of purely universally quantified formulas or purely existentially quantified formulas is undecidable.
- The theory of reals with addition and multiplication is decidable!
- The theory of rationals with addition and multiplication is undecidable.

- Checking whether a set of axioms has no contradiction is undecidable.

Chapter 1

Propositional Logic

Modern logic is a formal, symbolic system that tries to capture the principles of correct reasoning and truth. To describe any formal language precisely, we need three pieces of information — the *alphabet* describes the symbols used to write down the sentences in the language; the *syntax* describes the rules that must be followed for describing “grammatically correct” sentences in the language; and finally, the *semantics* gives “meaning” to the sentences in our formal language. You may have encountered other contexts where formal languages were introduced in such a manner. Here are some illustrative examples.

Example 1.1 Binomial coefficients are written using natural numbers and parentheses. However, not every way to put together parenthesis and natural numbers is a binomial coefficient. For example, $(1, 1)$, or $(^2)$ are examples of things that are no binomial coefficients. Correctly formed binomial coefficients are of the form $(^{i+j})_i$, where i and j are natural numbers. We could define the meaning of $(^{i+j})_i$ to be the natural number $\frac{(i+j)!}{i!j!}$. On the other hand, we could define the meaning of $(^{i+j})_i$ to be the number of ways of choosing i elements from a set of $i + j$ elements. Though both these ways of interpreting binomial coefficients are the same, they have a very different presentation. In general, one could define semantics in different ways, or even very different semantics to the same syntactic objects.

Example 1.2 Precise definitions of programming languages often involve characterizing its syntax and semantics. Turtle is an extremely simple programming language for drawing pictures. Programs in this language are written using **F**, **+**, and **-**. Any sequence formed by such symbols is a syntactically correct program in this language. We will interpret such a sequence of symbols as instructions to draw a picture — **F** is an instruction to draw a line by moving forward 1 unit; **+** is an instruction to turn the heading direction 60° to the left; **-** is an instruction to turn the heading direction 60° to the right. Figure 1.1 shows example programs and the pictures they draw based on this interpretation.

Even though Turtle is a very simple programming language, some very interesting curves can be approximated. Consider the following iterative procedure that produces

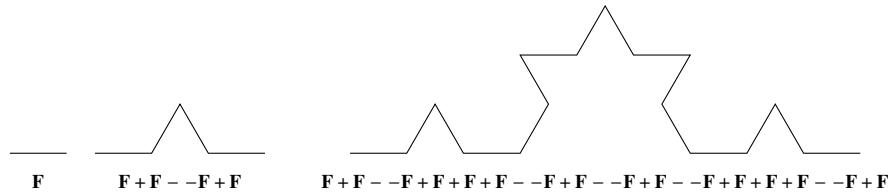


Fig. 1.1 Example Turtle programs and the pictures they draw.

a sequence of programs. Start with the program **F**. In each iteration, if P is a program at the start of the iteration, then construct the program P' obtained by replacing each **F** in P by **F + F - -F + F**. So at the beginning we have program **F**, in the next iteration the program is **F + F - -F + F**, and in the iteration after that it will be **F + F - -F + F + F - -F + F - -F + F**, and so on. The programs in this sequence draw pictures that in the limit approach the Koch curve.

Example 1.3 Regular expressions define special collections of strings called regular languages. Regular expressions over an alphabet Σ are built up using Σ , parentheses, \emptyset , ε , $,$, $+$, and $*$. Inductively, they are defined as the smallest set that satisfy the following rules.

- \emptyset and ε are regular expressions.
- For any $a \in \Sigma$, a is a regular expression.
- If r_1, r_2 are regular expressions then so are $(r_1 \cdot r_2)$, $(r_1 + r_2)$, and (r_1^*) .

Each regular expression r , semantically defines a subset of Σ^* ¹ that we will denote by $\llbracket r \rrbracket$. The semantics of regular expressions is defined inductively as follows.

- $\llbracket \emptyset \rrbracket = \emptyset$, and $\llbracket \varepsilon \rrbracket = \{\varepsilon\}$.
- For $a \in \Sigma$, $\llbracket a \rrbracket = \{a\}$.
- Inductively, $\llbracket (r_1 + r_2) \rrbracket = \llbracket r_1 \rrbracket \cup \llbracket r_2 \rrbracket$, $\llbracket (r_1 \cdot r_2) \rrbracket = \llbracket r_1 \rrbracket \cdot \llbracket r_2 \rrbracket$ and $\llbracket (r_1^*) \rrbracket = \llbracket r_1 \rrbracket^*$, where \cdot (on the right hand side) denotes the concatenation of two languages, and $*$ denotes the Kleene closure of a language.

We will now define one of the simplest logics encountered in an introductory discrete mathematics class called *propositional* or *sentential* logic. Propositional logic is a symbolic language to reason about *propositions*. Propositions are declarative sentences that are either true or false. Examples of such include “Springfield is the capital of Illinois”, “ $1+1 = 2$ ”, “ $2+2 = 0$ ”. Notice that propositions don’t need to be true facts and their truth may depend on the context. For example, “ $2+2 = 0$ ” is not true under the standard interpretation of $+$ as integer addition, but is true if $+$ denotes addition modulo 4. Non-examples of propositions include questions (like “What is it?”), commands (like “Read this!”), and things like “Location of robot”.

¹ For a finite set Σ , Σ^* denotes the collection of (finite) sequences/strings/words over Σ . For $n \in \mathbb{N}$, we use Σ^n to denote the set of sequences/strings/words over Σ of length exactly n .

The logic itself will be symbolic and abstract away from English sentences like the ones above. We will introduce a precise definition of this logic, much in the same way as Example 1.3, defining the syntax and semantics inductively.

1.1 Syntax

We will assume a (countably infinite) set of *propositions* $\text{Prop} = \{p_i \mid i \in \mathbb{N}\}$. The formulas of propositional logic will be strings over the alphabet $\text{Prop} \cup \{(\cdot), \neg, \vee\}$. Here \neg is *negation*, and \vee is *disjunction*.

Definition 1.4 The set of *well formed formulas* (wff) in propositional logic (over the set Prop) is the smallest set satisfying the following properties.

1. Any proposition $p_i \in \text{Prop}$ is a wff.
2. If φ is a wff then $(\neg\varphi)$ is a wff.
3. If φ and ψ are wffs then $(\varphi \vee \psi)$ is a wff.

Examples of wffs include p_1 , $(\neg p_1)$, $(p_1 \vee p_2)$, $((\neg(p_1 \vee p_3)) \vee (p_1 \vee p_4))$. On the other hand the following strings are not wffs: $(p_1 \neg)$, $(\vee p_1)$, $(p_1 \vee)$.

Inductive definitions of the kind in Example 1.3 or Definition 1.4 are quite common when defining the syntax of formulas in a logic or of programming languages. Therefore, in computer science, one often uses a “grammar-like” presentation for the syntax. For example, wffs φ in propositional logic are given by the following *BNF grammar*.

$$\varphi ::= p \mid (\neg\varphi) \mid (\varphi \vee \varphi) \quad (1.1)$$

where p is an element of Prop . Reading such grammars takes some getting used to. For example, the rule $(\varphi \vee \varphi)$ doesn’t mean that disjunctions can only be used when the two arguments are the same. Instead it says that if we take two elements that belong to the syntactic entity φ (i.e., wffs), put \vee between them with surrounding parenthesis, then we get another element belonging to the same syntactic entity φ . We will sometimes use such a grammar representation to describe syntax in a succinct manner.

Inductive Definitions

What is the set identified by inductive definitions like Definition 1.4 and (1.1) ? Is there a unique single minimal set that satisfies the conditions in Definition 1.4? After all sets can be incomparable with respect to the \subseteq relation. And what does the grammar given in (1.1) mean? Finally, do Definition 1.4 and (1.1) identify the same set?

Let us begin by first defining the set described by (1.1) . Equation (1.1) defines the set $S = \bigcup_{i \in \mathbb{N}} S_i$, where the sets S_i (for $i \in \mathbb{N}$) are given as follows.

$$\begin{aligned} S_0 &= \text{Prop} \\ S_{i+1} &= S_i \cup \{(\neg\varphi) \mid \varphi \in S_i\} \cup \{(\psi \vee \varphi) \mid \psi, \varphi \in S_i\} \quad \text{for } i \geq 0 \end{aligned}$$

Note that $S = \bigcup_{i \in \mathbb{N}} S_i$ is an infinite union. You can think of S as the limit of the increasing sequence of sets $S_0 \cup S_1 \cup \dots \cup S_n$. Another way to think of S is as the set of all φ that belong to *some* S_i . That is,

$$S = \{\varphi \mid \varphi \in S_i, \text{ for some } i\}.$$

Another way to interpret the sets S and S_i is as follows. S_i denotes the set of formulas that can be derived from the grammar rules within i steps, and S is the set of formulas that can be derived from the grammar rules in *some finite number* of steps.

Given the above meaning, it's natural to prove properties about the set of expressions S using induction on i . More precisely, if we want to show a property P is true about S , then we:

- Establish P to be true for every expression in S_0 .
- For every $i \geq 0$, we assume that P holds for every expression in S_i and prove it holds for all expressions in S_{i+1} .

The above shows that P holds on S_i , for every $i \in \mathbb{N}$, and hence P holds for every formula in S .

Let us now consider Definition 1.4 and argue that it is well defined. Before showing that there is a unique smallest set satisfying conditions (1), (2), and (3) of Definition 1.4, let us argue that there is at least one set satisfying (1), (2), and (3). Take the set of all possible strings built over the alphabet $\text{Prop} \cup \{(\cdot), \neg, \vee\}$. This set clearly satisfies all the conditions. But why should there be a smallest set? Observe that if two sets A and B satisfy the conditions in Definition 1.4, then so does $A \cap B$. More generally, if $\{A_i\}_{i \in I}$ is a (possibly infinite) collection of sets satisfying (1), (2), and (3), then so does the set $\bigcap_{i \in I} A_i$. Thus, the intersection of *all* sets which satisfy the conditions, is the unique, smallest (with respect to set inclusion) set identified by Definition 1.4. Hence, it is well defined.

Let S denote the set identified by grammar in (1.1) and T the set defined in Definition 1.4. We will argue that these two sets are the same. First, let us show that S satisfies conditions (1), (2), and (3) of Definition 1.4. Observe that, by definition, for every i , $S_i \subseteq S_{i+1}$. Hence, the sets increase with index, i.e., for every $i < j$, $S_i \subseteq S_j$; this can be formally established by induction, and we leave this as an exercise. Since $\text{Prop} \subseteq S_0$, this means $\text{Prop} \subseteq S_i$ for every i , and therefore S satisfies condition (1). Next, consider any $\varphi, \psi \in S$. By definition of S , this means there is are $i, j \in \mathbb{N}$ such that $\varphi \in S_i$ and $\psi \in S_j$. Taking $k = \max(i, j)$, we can conclude that $\{\varphi, \psi\} \subseteq S_k$. Thus, by definition, $(\neg\varphi)$ and $(\varphi \vee \psi)$ both belong to S_{k+1} , and hence are in S . Therefore, S satisfies conditions (2) and (3) of Definition 1.4. Finally, since S satisfies all conditions of Definition 1.4 and T is the smallest set with these properties, we can conclude that $T \subseteq S$.

To prove the other inclusion (that $S \subseteq T$), we will prove that for every i , $S_i \subseteq T$ by induction. In the base case observe that $S_0 = \text{Prop} \subseteq T$, since T satisfies condition (1). Assume as induction hypothesis, that for all $j \leq i$, $S_j \subseteq T$. In the induction step, we need to establish the claim that $S_{i+1} \subseteq T$. Let $\varphi \in S_{i+1}$ be an arbitrary element. By the definition of S_{i+1} , there are 3 cases to consider. If $\varphi \in S_i$ then by induction hypothesis $\varphi \in T$. If $\varphi = (\neg\psi)$ for some $\psi \in S_i$, then by induction hypothesis $\psi \in T$,

and since T satisfies condition (2), $\varphi = (\neg\psi) \in T$. Finally, if $\varphi = (\psi_1 \vee \psi_2)$ for some $\psi_1, \psi_2 \in S$; then by induction hypothesis, $\psi_1, \psi_2 \in T$ and since T satisfies condition (3), $\varphi = (\psi_1 \vee \psi_2) \in T$. This completes the proof that $S \subseteq T$, and therefore, $S = T$.

The above argument, establishing that the smallest set satisfying some closure conditions is the same as the set of objects derived from grammar rules, can be generalized to any grammar (not just context-free grammars). In general, if one defines a set as the smallest set S that satisfies conditions of the form “if these elements belong to S than these other elements must belong to S ”, then the smallest set is well-defined. But if you also have conditions saying “if these elements belong to S , these elements should *not* belong to S ,” then the “smallest” set may not be a well-defined.

Inductive or *recursive* definitions are ubiquitous in computer science, and reasoning about such definitions is naturally done using some form of induction. In fact, the prevalence of induction in computer science is because of the ubiquity of recursive definitions. Here are some other examples of recursive definitions.

- Consider the operational semantics of a program. The set of states/configurations that the program can reach is best defined recursively. It is the smallest set R of states such that (a) R contains the initial states of the program, and (b) if a state s is in R and the program can transition in one step from s to s' , then $s' \in R$. Because of the recursive nature of this definition, reasoning about programs often involves induction. For example, to show that a program does not throw an exception, one needs to prove that any reachable state is one that does not throw an exception. This is typically established using induction.
- Consider *lists* in programs. The `cons` operator constructs a new list by adding an element to the head of another list. Lists over the elements E can then be defined as the smallest set L such that (a) L contains `nil`, the empty list, and (b) if $\ell \in L$ and $e \in E$, then `cons`(e, ℓ) is in L as well.
- The set of *natural numbers* can also be defined recursively. Let us denote by `succ`, the successor function ². Then the set of natural numbers is the smallest set \mathbb{N} such that (a) \mathbb{N} contains 0, and (b) for every $n \in \mathbb{N}$, `succ`(n) belongs to \mathbb{N} .

Other logical operators and Operator Precedence

Conjunction and *implication* are logical operators that arise quite often when expressing properties. These operators can be defined in terms of \neg and \vee . Let φ and ψ be wffs. $(\varphi \wedge \psi)$ (read “ φ and ψ ”) denotes the formula $(\neg((\neg\varphi) \vee (\neg\psi)))$. And $(\varphi \rightarrow \psi)$ (read “ φ implies ψ ”) denotes the formula $((\neg\varphi) \vee \psi)$. Another useful wff is \top (read “true”); it denotes the formula $(p \vee (\neg p))$, where p is (any) proposition. Finally the wff \perp (read “false”) is the formula $(\neg\top)$.

² Intuitively, given a number n , `succ`(n) is the next number, namely, $n + 1$. However, this interpretation is only in our minds. `succ` is just some function.

Writing formulas strictly according to the syntax presented is cumbersome because of many parentheses and subscripts. Therefore, we will make the following notational simplifications.

- The outermost parentheses will be dropped. Thus we will write $p_3 \vee (p_2 \vee p_1)$ instead of $(p_3 \vee (p_2 \vee p_1))$
- We will sometimes omit subscripts of propositions. Thus we will write p instead of p_1 , or q instead of p_2 , r instead of p_3 , or s instead of p_4 , and so on.
- The following precedence of operators will be assumed: \neg , \wedge , \vee , \rightarrow . Thus $\neg p \wedge q \rightarrow r$ will mean $((\neg p) \wedge q) \rightarrow r$.

Definition 1.4 for wffs in propositional logic has the nice property that the structure of a formula can be interpreted in a unique way. There is no ambiguity in its interpretation. For example, if $\neg p_1 \vee p_2$ were a wff, then it is unclear whether we mean the formula $\varphi = ((\neg p_1) \vee p_2)$ or $\psi = (\neg(p_1 \vee p_2))$ ³ — in $\varphi \vee$ is the topmost operator, while in $\psi \neg$ is the topmost operator. Our syntax does not have such issues. This will be exploited often in inductive definitions and in algorithms. This observation can be proved by structural induction, but we skip its proof.

Theorem 1.5 (Unique Readability)

Any wff can be uniquely read, i.e., it has a unique topmost logical operator and well defined immediate sub-formulas.

1.2 Semantics

We will now provide a meaning or *semantics* to the formulas. Our definition will follow the inductive definition of the syntax, just like in Example 1.3. The semantics of formulas in a logic, are typically defined with respect to a *model*, which identifies a “world” in which certain facts are true. In the case of propositional logic, this world or model is a *truth valuation* or *assignment* that assigns a truth value (true/false) to every proposition. The *truth value* *truth* will be denoted by T , and the truth value *falsity* will be denoted by F .

Definition 1.6 A (*truth*) *valuation* or *assignment* is a function v that assigns truth values to each of the propositions, i.e., $v : \text{Prop} \rightarrow \{T, F\}$.

The value of a proposition p under valuation v is given by $v(p)$.

We will define the semantics through a *satisfaction relation*, which is a binary relation \models between valuations and formulas. The statement $v \models \varphi$ should be read as “ v satisfies φ ” or “ φ is true in v ” or “ v is a model of φ ”. It is defined inductively following the syntax of formulas. In the definition below, we say $v \not\models \varphi$ when $v \models \varphi$ does not hold.

Definition 1.7 For a valuation v and wff φ , the satisfaction relation, $v \models \varphi$, is defined inductively based on the structure of φ as follows.

³ For the formulas here, we are not using the precedence rules given before.

- $v \models p$ if and only if $v(p) = T$.
- $v \models (\neg\varphi)$ if and only if $v \not\models \varphi$.
- $v \models (\varphi \vee \psi)$ if either $v \models \varphi$ or $v \models \psi$.

Example 1.8 Let us look at a couple of examples to see how the inductive definition of the satisfaction relation can be applied. Consider the formula $\varphi = \neg(\neg p \vee \neg q) \vee (\neg p \vee \neg q)$. Recall with respect to the notational simplifications we identified, φ is the formula $((\neg((\neg p) \vee (\neg q))) \vee ((\neg p) \vee (\neg q)))$. Consider the valuation v_1 that sets all propositions to T . Now $v_1 \models \varphi$ can be seen from the following observations.

| | |
|--|----------------------|
| $v_1 \models p$ | because $v_1(p) = T$ |
| $v_1 \not\models \neg p$ | semantics of \neg |
| $v_1 \models q$ | because $v_1(q) = T$ |
| $v_1 \not\models \neg q$ | semantics of \neg |
| $v_1 \not\models \neg p \vee \neg q$ | semantics of \vee |
| $v_1 \models \neg(\neg p \vee \neg q)$ | semantics of \neg |
| $v_1 \models \neg(\neg p \vee \neg q) \vee (\neg p \vee \neg q)$ | semantics of \vee |

Consider v_2 that assigns all propositions to F . Once again $v_2 \models \varphi$. The reasoning behind this observation is as follows.

| | |
|--|----------------------|
| $v_2 \not\models p$ | because $v_2(p) = F$ |
| $v_2 \models \neg p$ | semantics of \neg |
| $v_2 \models \neg p \vee \neg q$ | semantics of \vee |
| $v_2 \models \neg(\neg p \vee \neg q) \vee (\neg p \vee \neg q)$ | semantics of \vee |

The semantics in Definition 1.7 defines a satisfaction relation between valuations and formulas. However, one could define the semantics of propositional logic differently, by considering the formula as a “program” or “circuit” that computes a truth value based on the assignment. This approach is captured by the following definition of the *value* of a wff under a valuation.

Definition 1.9 The *value* of a wff φ under valuation v , denoted by $v[\varphi]$, is inductively defined as follows.

$$\begin{aligned} v[p] &= v(p) \\ v[(\neg\varphi)] &= \begin{cases} F & \text{if } v[\varphi] = T \\ T & \text{if } v[\varphi] = F \end{cases} \\ v[\varphi \vee \psi] &= \begin{cases} F & \text{if } v[\varphi] = v[\psi] = F \\ T & \text{otherwise} \end{cases} \end{aligned}$$

Example 1.10 Let us consider $\varphi = \neg(\neg p \vee \neg q) \vee (\neg p \vee \neg q)$ and v_1 which assigns all propositions to T . $v_1[\varphi]$ can be computed as follows.

| | |
|---|----------------------|
| $v_1[p] = T$ | because $v_1(p) = T$ |
| $v_1[\neg p] = F$ | semantics of \neg |
| $v_1[q] = T$ | because $v_1(q) = T$ |
| $v_1[\neg q] = F$ | semantics of \neg |
| $v_1[\neg p \vee \neg q] = F$ | semantics of \vee |
| $v_1[\neg(\neg p \vee \neg q)] = T$ | semantics of \neg |
| $v_1[\neg(\neg p \vee \neg q) \vee (\neg p \vee \neg q)] = T$ | semantics of \vee |

Definitions 1.7 and 1.9 are both equivalent in some sense. This is captured by the following theorem.

Theorem 1.11 *For any truth valuation v and wff φ , $v \models \varphi$ if and only if $v[\varphi] = T$*

The proof of Theorem 1.11 is by structural induction on the formula φ . It is left as an exercise for the reader.

It is convenient to associate with every wff the set of truth valuations under which the formula holds.

Definition 1.12 The *models* of wff φ is the set of valuations that *satisfy* φ . More precisely,

$$[\varphi] = \{v \mid v \models \varphi\}.$$

Observe that as per the definition, $[\perp] = \emptyset$.

The *relevance lemma* says that whether φ holds under a valuation depends only on how the valuation maps the propositions that *syntactically occur* in the formula. This is intuitively obvious; surely, whether $(p \wedge q) \vee r$ holds is independent of whether the proposition s is mapped to true/false. For a wff φ , the set of propositions appearing in φ , denoted $\text{occ}(\varphi)$, is inductively defined as follows.

$$\begin{aligned}\text{occ}(p) &= \{p\} \\ \text{occ}(\neg\varphi) &= \text{occ}(\varphi) \\ \text{occ}(\varphi \vee \psi) &= \text{occ}(\varphi) \cup \text{occ}(\psi)\end{aligned}$$

The relevance lemma is then as follows.

Lemma 1.13 (Relevance Lemma)

Let v_1 and v_2 be truth valuations such that for all $p \in \text{occ}(\varphi)$, we have $v_1(p) = v_2(p)$, i.e., v_1 and v_2 agree on the truth values assigned to all propositions in $\text{occ}(\varphi)$. Then $v_1 \models \varphi$ if and only if $v_2 \models \varphi$.

Proof By structural induction on φ .

Base Case $\varphi = p$ Observe that, $v_1 \models \varphi$ iff $v_1(p) = T = v_2(p)$ iff $v_2 \models \varphi$.

Induction Step $\varphi = (\neg\psi)$ Since $\text{occ}(\psi) \subseteq \text{occ}(\varphi)$, we have by induction hypothesis, $v_1 \models \psi$ iff $v_2 \models \psi$. Therefore, by the semantics of \neg , $v_1 \models \varphi$ iff $v_2 \models \varphi$.

Induction Step $\varphi = (\psi_1 \vee \psi_2)$ Since $\text{prop}(\psi_i) \subseteq \text{prop}(\varphi)$ (for $i \in \{1, 2\}$), we have by induction hypothesis, $v_1 \models \psi_i$ iff $v_2 \models \psi_i$. Therefore, by the semantics of \vee , $v_1 \models \varphi$ iff $v_2 \models \varphi$. \square

Lemma 1.13 implies that, to determine if a formula holds in a valuation, we only need to consider the assignment to the finitely many propositions occurring in the formula. Thus, instead of thinking of valuations as assigning truth values to all (infinitely many) propositions, we can think of them as functions with a finite domain.

1.3 Satisfiability and Validity

Two formulas that are syntactically different, could however, be “semantically equivalent”. But what do we mean by semantic equivalence? Intuitively, this is when the truth value of each formula in every valuation is the same.

Definition 1.14 (Logical Equivalence)

A wff φ is said to be *logically equivalent* to ψ iff any of the following equivalent conditions hold.

- for every valuation v , $v \models \varphi$ iff $v \models \psi$,
- for every valuation v , $v[\varphi] = v[\psi]$,
- $[\varphi] = [\psi]$.

We denote this by $\varphi \equiv \psi$.

Let us consider an example to see how we may reason about two formulas being logically equivalent.

Example 1.15 Consider the wffs $\varphi_1 = p \wedge (q \vee r)$ and $\varphi_2 = (p \wedge q) \vee (p \wedge r)$, where p , q and r are propositions. Though φ_1 and φ_2 are syntactically different, they are semantically equivalent. To prove that $\varphi_1 \equiv \varphi_2$, we need to show that they two formulas evaluate to the same truth value under every valuation. One convenient way to organize such a proof is as *truth table*, where different cases in the case-by-case analysis correspond to different rows. Each row of the truth table corresponds to a (infinite) collection of valuations based on the value assigned to propositions p , q and r ; the columns correspond to the value of different (sub)-formulas under each valuation in this collection. For example, a truth table reasoning for φ_1 and φ_2 will look as follows.

Notice that since the columns corresponding to φ_1 and φ_2 are identical in every row, and every valuation corresponds to some row in the table, it follows that φ_1 and φ_2 are logically equivalent.

Let us now consider $\varphi'_1 = \psi_1 \wedge (\psi_2 \vee \psi_3)$ and $\varphi'_2 = (\psi_1 \wedge \psi_2) \vee (\psi_1 \wedge \psi_3)$, where ψ_1, ψ_2 and ψ_3 are arbitrary wffs. Once again φ'_1 and φ'_2 are logically equivalent, no matter what ψ_1, ψ_2 and ψ_3 are. The reasoning is essentially the same as above. The rows of the truth table now classify valuations based on the value of formulas ψ_1, ψ_2 and ψ_3 under them.

| ψ_1 | ψ_2 | ψ_3 | $\psi_2 \vee \psi_3$ | φ_1 | $\psi_1 \wedge \psi_2$ | $\psi_1 \wedge \psi_3$ | φ_2 |
|----------|----------|----------|----------------------|-------------|------------------------|------------------------|-------------|
| F | F | F | F | F | F | F | F |
| F | F | T | T | F | F | F | F |
| F | T | F | T | F | F | F | F |
| F | T | T | T | F | F | F | F |
| T | F | F | F | F | F | F | F |
| T | F | T | T | T | F | T | T |
| T | T | F | T | T | T | F | T |
| T | T | T | T | T | T | T | T |

Truth table based reasoning, as carried out in Example 1.15, is a very convenient way to organize proofs of propositional logic. We will often use it. Example 1.15 highlights another important observation. Let φ and ψ be logically equivalent formulas. Let φ' and ψ' be formulas obtained by substituting propositions occurring in φ and ψ by arbitrary formulas. Then $\varphi' \equiv \psi'$.

Definition 1.16 (Logical Consequence)

Let Γ be a (possibly infinite) set of formulas and let φ be a wff. We say that φ is a *logical consequence* of Γ (denoted $\Gamma \models \varphi$) iff for every valuation v , if for every $\psi \in \Gamma, v \models \psi$ then $v \models \varphi$. In other words, any model that satisfies every formula in Γ also satisfies φ .

We could equivalently have defined it as $\Gamma \models \varphi$ iff $\bigcap_{\psi \in \Gamma} [\psi] \subseteq [\varphi]$.

Example 1.17 Consider the set $\Gamma = \{\psi_1 \rightarrow \psi_2, \psi_2 \rightarrow \psi_1, \psi_1 \vee \psi_2\}$, where ψ_1 and ψ_2 are arbitrary formulas. We will show that $\Gamma \models \psi_1$. Once again, we will use a truth table to classify valuations into row based on the value that ψ_1 and ψ_2 evaluate to. Such a truth table looks as follows.

| ψ_1 | ψ_2 | $\psi_1 \rightarrow \psi_2$ | $\psi_2 \rightarrow \psi_1$ | $\psi_1 \vee \psi_2$ |
|----------|----------|-----------------------------|-----------------------------|----------------------|
| F | F | T | T | F |
| F | T | T | F | T |
| T | F | F | T | T |
| T | T | T | T | T |

Notice that there is only one row where columns 3, 4, and 5 are all T; this corresponds to the valuations where ψ_1 and ψ_2 evaluate to T, and under every such valuation, all formulas in Γ are satisfied. In this row, since ψ_1 also evaluates to T, we have that $\Gamma \models \psi_1$.

It is worth observing one special case of Definition 1.16 — when $\Gamma = \emptyset$. In this case, every valuation satisfies every formula in Γ (vacuously, since there are none to satisfy). Therefore, if $\emptyset \models \varphi$, then every truth assignment satisfies φ . Such formulas are called *tautologies*, and they represent universal truths that hold in every model/world/assignment.

Definition 1.18 (Tautologies)

A wff φ is a *tautology* or is *valid* if for every valuation v , $v \models \varphi$. In other words, $\emptyset \models \varphi$. We will denote this simply as $\models \varphi$.

Example 1.19 We will show $\varphi = \psi_1 \rightarrow (\psi_2 \rightarrow \psi_1)$ is a tautology, no matter what formulas ψ_1 and ψ_2 are. The proof is once again organized as truth table, and we show that in all rows the formula φ evaluates to \top .

| ψ_1 | ψ_2 | $\psi_2 \rightarrow \psi_1$ | $\psi_1 \rightarrow (\psi_2 \rightarrow \psi_1)$ |
|----------|----------|-----------------------------|--|
| F | F | T | T |
| F | T | F | T |
| T | F | T | T |
| T | T | T | T |

The last important notion we would like to introduce is that of *satisfiability*.

Definition 1.20 (Satisfiability)

A formula φ is *satisfiable* if there is some valuation v such that $v \models \varphi$. In other words, $\llbracket \varphi \rrbracket \neq \emptyset$. If a formula is not satisfiable, we say it is *unsatisfiable*.

Example 1.21 $\varphi = (p \vee q) \wedge (\neg p \vee \neg q)$ is satisfiable because the valuation v that maps p to \top and q to \top satisfies φ , i.e., $v \models \varphi$.

Based on Definitions 1.18 and 1.20, it is easy to see that there is a close connection between satisfiability and validity.

Proposition 1.22 A wff φ is valid if and only if $\neg\varphi$ is unsatisfiable.

Proof Let v be any valuation. If φ is valid, we know that $v \models \varphi$. Therefore, by the semantics of \neg , we have $v \not\models \neg\varphi$. Thus $\neg\varphi$ is unsatisfiable. Conversely, if $\neg\varphi$ is unsatisfiable, then $v \not\models \neg\varphi$. Again, by the semantics of \neg , $v \models \varphi$. Thus, φ is valid. \square

We conclude this section by considering two fundamental computational problems — *satisfiability* and *validity*.

Satisfiability Given a formula φ , determine if φ is satisfiable.

Validity Given a formula φ , determine if φ is a tautology.

The satisfiability and validity problems have very simple algorithms to solve them. To check if φ , over propositions $\{p_1, \dots, p_n\}$, is satisfiable (is a tautology), compute $v \llbracket \varphi \rrbracket$ for every truth assignment v to the propositions $\{p_1, \dots, p_n\}$. The running time for this algorithm is $O(2^n)$. One of the most important open questions in computer science is whether this is the best algorithm for these problems. The following theorem by Cook and Levin, supports the belief that this exponential algorithm is unlikely to be improved in the worst case.

Theorem 1.23 (Cook-Levin)

The satisfiability problem for propositional logic is NP-complete.

Proof Any proof showing a problem to be NP-complete has two parts. First is an argument that the problem belongs to NP and the second that it is hard.

Membership in NP. Given a formula φ , the NP-algorithm to check satisfiability is as follows — Guess a truth assignment v , evaluate φ on the truth assignment, and accept if $v[\varphi] = T$; otherwise reject. Guessing (nondeterministically) a truth assignment takes time which is linear in the number of propositions in φ , which is linear in the size of φ , and computing $v[\varphi]$ also takes time that is linear in the size of φ , where the evaluation algorithm computes the value in a “bottom-up” fashion. Thus, the total running time is polynomial.

NP-hardness. Consider $A \in \text{NP}$. Let $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \triangleright)$ be a nondeterministic TM recognizing A in time n^ℓ , where Q is the set of control states, Σ is the input alphabet, Γ is the tape alphabet, δ is the transition function, q_0 is the initial state, q_{acc} is the unique accepting state, q_{rej} is the unique rejecting state, \sqcup is the blank symbol, and \triangleright is the left end marker symbol that appears on the leftmost cell of each tape. Without loss of generality, we assume that q_{acc} and q_{rej} are the only halting states of M . We will also assume that M has a read only input tape, and a *single* read/write work tape.

For an input x , we will construct (in polynomial time) a formula $f_M(x)$ such that M accepts x (i.e., $x \in A$) iff $f_M(x)$ is satisfiable. $f_M(x)$ will encode constraints on a computation of M on x such that a satisfying assignment to $f_M(x)$ will describe “how M accepts x ”. That is, $f_M(x)$ will encode that

- M starts in the initial configuration with input x ,
- Each configuration follows from the previous one in accordance with the transition function of M ,
- The accepting state is reached in the last step.

Let us formalize this intuition by giving a precise construction. We begin by identifying the set of propositions we will use, and their informal interpretation.

Propositional Variables. The propositions of $f_M(x)$ will be as follows.

| Name | Meaning if set to T | Total Number |
|-----------------------|--|-------------------------|
| InpSymb(b, p) | Input tape stores b at position p | $O(x)$ |
| TapeSymb(b, p, i) | Work tape stores b in cell p at time i | $O(x ^{2\ell})$ |
| InpHd(h, i) | Input head in cell h at time i | $O(x \cdot x ^\ell)$ |
| TapeHd(h, i) | Work tape’s head in cell h at time i | $O(x ^{2\ell})$ |
| State(q, i) | State is q at time i | $O(x ^\ell)$ |

Abbreviations. In order to define $f_M(x)$, the following abbreviations will be useful.

- $\bigwedge_{k=1}^m X_k$ means $X_1 \wedge X_2 \wedge \dots \wedge X_m$

- $\nabla(X_1, X_2, \dots, X_m)$ will denote a formula that is satisfiable iff exactly one of X_1, \dots, X_m is set to true. In other words,

$$\nabla(X_1, X_2, \dots, X_m) = (X_1 \vee X_2 \vee \dots \vee X_m) \wedge \bigwedge_{k \neq l} (\neg X_k \vee \neg X_l)$$

Overall Reduction. The overall form of $f_M(x)$ will be as follows.

$$f_M(x) = \varphi_{\text{initial}} \wedge \varphi_{\text{consistent}} \wedge \varphi_{\text{transition}} \wedge \varphi_{\text{accept}}$$

where

- φ_{initial} says that “configuration at time 0 is the initial configuration with input x ”
- $\varphi_{\text{consistent}}$ says that “at each time, truth values to variables encode a valid configuration”
- $\varphi_{\text{transition}}$ says that “configuration at each time follows from the previous one by taking a transition”
- φ_{accept} says that “the last configuration is an accepting configuration”

We now outline what each of the above formulas is.

Initial Conditions Let $x = a_1 a_2 \dots a_n$

$$\begin{aligned} \varphi_{\text{initial}} = & \text{State}(q_0, 0) && \text{“At time 0, state is } q_0\text{”} \\ & \wedge \text{InpSymb}(\triangleright, 0) \wedge \text{TapeSymb}(\triangleright, 0, 0) && \text{“Leftmost cells contain } \triangleright\text{”} \\ & \wedge \bigwedge_{p=1}^n \text{InpSymb}(a_p, p) && \text{“At time 0, cells 1 through } n\text{ hold } x\text{”} \\ & \wedge \bigwedge_{p=1}^{n^\ell} \text{TapeSymb}(\square, p, 0) && \text{“At time 0, all work tape cells are blank”} \\ & \wedge \text{InpHd}(0, 0) \wedge \text{TapeHd}(0, 0) && \text{“At time 0, all heads at the leftmost position”} \end{aligned}$$

Consistency Assume that the tape alphabet is $\Gamma = \{b_1, b_2, \dots, b_t\}$ and the set of states is $Q = \{q_0, q_1, \dots, q_m\}$.

$$\begin{aligned}
\varphi_{\text{consistent}} = & \bigwedge_{i=0}^{n^\ell} \nabla(\text{State}(q_0, i), \dots, \text{State}(q_m, i)) \\
& \quad \text{"At any time } i, \text{ state is unique"} \\
& \bigwedge_{i=0}^{n^\ell} \text{TapeSymb}(\triangleright, 0, i) \\
& \quad \text{"At any time, leftmost cell contains } \triangleright" \\
& \bigwedge_{i=0}^{n^\ell} \bigwedge_{p=0}^{n^\ell} \nabla(\text{TapeSymb}(b_1, p, i), \dots, \text{TapeSymb}(b_t, p, i)) \\
& \quad \text{"work tape cells contain unique symbols"} \\
& \bigwedge_{i=0}^{n^\ell} \nabla(\text{InpHd}(0, i), \dots, \text{InpHd}(n, i)) \\
& \quad \text{"At any time, input head is in one cell"} \\
& \bigwedge_{i=0}^{n^\ell} \nabla(\text{TapeHd}(0, i), \dots, \text{TapeHd}(n^\ell, i)) \\
& \quad \text{"At any time, work tape head is in one cell"}
\end{aligned}$$

Transition Consistency Consider a non-halting state q (i.e., $q \notin \{q_{\text{acc}}, q_{\text{rej}}\}$), input symbol c_{in} and tape symbol c_w . Let the transition at state q , when reading these symbols be given by

$$\delta(q, c_{in}, c_w) = \{(q^{(1)}, d_{in}^{(1)}, c_w^{(1)}, d_w^{(1)}), \dots, (q^{(s)}, d_{in}^{(s)}, c_w^{(s)}, d_w^{(s)})\}.$$

Here $(q^{(i)}, d_{in}^{(i)}, c_w^{(i)}, d_w^{(i)}) \in \delta(q, c_{in}, c_w)$ means that if M is in state q and reads c_{in} on the input tape and c_w on the work tape, then one possible transition is to state $q^{(i)}$, moving the input head in direction $d_{in}^{(i)}$, writing $c_w^{(i)}$ on the work tape, and moving the work tape head in direction $d_w^{(i)}$. Direction -1 denotes moving the head *left* and $+1$ denotes moving the head *right*. We will first define a formula $\Delta_{q, c_{in}, c_w}^{i, p_{in}, p_w}$ that says that at time i if the state is q and the symbol read on the input tape is c_{in} and on the work tape is c_w at positions p_{in}, p_w , respectively, then at time $i + 1$ the state, symbols written and new head position is one of the tuples described by the δ function.

$$\begin{aligned}
\Delta_{q, c_{in}, c_w}^{i, p_{in}, p_w} = & (\text{State}(q, i) \wedge \text{InpHd}(p_{in}, i) \wedge \text{InpSymb}(c_{in}, p_{in}, i) \wedge \\
& \text{TapeHd}(p_w, i) \wedge \text{TapeSymb}(c_w, p_w, i)) \rightarrow \\
& \nabla(\text{Ch}_{i, p_{in}, p_w}^1, \text{Ch}_{i, p_{in}, p_w}^2, \dots, \text{Ch}_{i, p_{in}, p_w}^s)
\end{aligned}$$

where

$$\begin{aligned}
\text{Ch}_{i, p_{in}, p_w}^t = & \text{State}(q^{(t)}, i + 1) \wedge \text{InpHd}(p_{in} + d_{in}^{(t)}, i + 1) \wedge \\
& \text{TapeSymb}(c_w^{(t)}, p_w, i + 1) \wedge \text{TapeHd}(p_w + d_w^{(t)}, i + 1)
\end{aligned}$$

For a halting state q (i.e., $q \in \{q_{\text{acc}}, q_{\text{rej}}\}$), we define $\Delta_{q, c_{in}, c_w}^{i, p_{in}, p_w}$ as saying that the state, symbols, and head positions don't change. In other words,

$$\begin{aligned}
\Delta_{q, c_{in}, c_w}^{i, p_{in}, p_w} = & (\text{State}(q, i) \wedge \text{InpHd}(p_{in}, i) \wedge \text{InpSymb}(c_{in}, p_{in}) \wedge \\
& \text{TapeHd}(p_w, i) \wedge \text{TapeSymb}(c_w, p_w, i)) \rightarrow \\
& (\text{State}(q, i + 1) \wedge \text{InpHd}(p_{in}, i + 1) \\
& \wedge \text{TapeSymb}(c_w, p_w, i + 1) \wedge \text{TapeHd}(p_w, i + 1))
\end{aligned}$$

when q is a halting state.

Now Transition Consistency itself can be defined as follows.

$$\varphi_{\text{transition}} = \bigwedge_{i=0}^{n^\ell} \bigwedge_{p_{in}=0}^n \bigwedge_{p_w=0}^{n^\ell} \{$$

$$\begin{aligned} & \bigwedge_{b \neq c} \neg \text{TapeHd}(p_w, i) \rightarrow \\ & \quad \neg(\text{TapeSymb}(b, p_w, i) \wedge \text{TapeSymb}(c, p_w, i+1)) \\ & \quad \text{“If head is not in some position, then symbol does not change”} \end{aligned}$$

$$\bigwedge_{q=q_0}^{q_m} \bigwedge_{c_{in}=b_1}^{b_t} \bigwedge_{c_w=b_1}^{b_t} \Delta_{q, c_{in}, c_w}^{i, p_{in}, p_w} \}$$

$$\text{“If head is in some position, then a transition is taken”}$$

Acceptance

$$\varphi_{\text{accept}} = \text{State}(q_{\text{acc}}, n^\ell)$$

We can argue that M accepts x if and only if $f_M(x)$ is satisfiable. Further $f_M(x)$ can be constructed in time that is polynomial in the size of x ; the size of M also plays a role but that is fixed. \square

The Cook-Levin Theorem (Theorem 1.23) is an important result in computer science. It was the first result establishing the intractability of a problem. Moreover, it also implies the intractability of the validity problem. This is because there is a formula φ is valid if and only if $\neg\varphi$ is unsatisfiable.

Proposition 1.24 *A formula φ is valid if and only if $\neg\varphi$ is unsatisfiable.*

Proof The proposition can be established by the following sequence of observations. φ is valid iff for every valuation v , $v[\varphi] = \top$ (definition of validity) iff for every valuation v , $v[\neg\varphi] = \perp$ (from the semantics of \neg) iff $\neg\varphi$ is unsatisfiable (definition of unsatisfiability). \square

Using Proposition 1.24 we establish the coNP -hardness of validity.

Theorem 1.25 *The validity problem for propositional logic is coNP -complete.*

Proof Observe that there is a simple NP algorithm to check that a formula φ is *not* valid — Guess a valuation v , and check that $v[\varphi] = \perp$. Since checking non-validity has a NP algorithm, it means that the validity problem has a coNP algorithm, and is therefore, in coNP .

To prove the coNP -hardness of the validity problem, we make the following observations. First, for any two problems A and B , if $A \leq_P B$ then $\overline{A} \leq_P \overline{B}$. Thus, from the NP -hardness of the satisfiability problem (Theorem 1.23), we can conclude that the problem of checking if a formula is *unsatisfiable* is coNP -hard. Since unsatisfiability is coNP -hard, it follows that validity is also coNP -hard based on the observation in Proposition 1.24. \square

1.4 Compactness Theorem

The compactness theorem is an important property about propositional logic. In this section, we will look at a couple of different proofs of this theorem.

A (finite or infinite) set of formulas Γ is *satisfiable* if there is a valuation v such that for every $\varphi \in \Gamma$, $v \models \varphi$ (or $v[\varphi] = T$); we will denote this by $v \models \Gamma$. A set of formulas Γ is *finitely satisfiable* if every finite subset Γ_0 of Γ is satisfiable. These two notions, satisfiability and finite satisfiability, are equivalent — this is the content of the *compactness theorem*.

Theorem 1.26 (Compactness)

A set of formulas Γ is satisfiable if and only if Γ is finitely satisfiable.

Observe that if Γ is satisfiable then the satisfying assignment (say v) also satisfies every subset of Γ and therefore also every finite subset of Γ . Thus if Γ is satisfiable then it is also finitely satisfiable. The challenge is, therefore, in proving the converse — that finite satisfiability implies satisfiability. If Γ is a finite set, then clearly finite satisfiability implies satisfiability because Γ itself is a finite subset of Γ . So the interesting case is when Γ is infinite. We will provide a couple of very different proofs for this result.

Before moving on to present our proofs for Theorem 1.26, we highlight an important consequence of the theorem.

Corollary 1.27 *Let Γ be a (possibly infinite) set of formulas and let φ be a formula. If $\Gamma \models \varphi$ then there is a finite subset $\Delta \subseteq \Gamma$ such that $\Delta \models \varphi$.*

Proof Let $\Gamma \models \varphi$. Then $\Gamma \cup \{\neg\varphi\}$ is unsatisfiable. By Theorem 1.26, there is a finite subset $\Delta' \subseteq \Gamma \cup \{\neg\varphi\}$ that is unsatisfiable. Taking $\Delta = \Delta' \setminus \{\neg\varphi\}$, we observe that $\Delta \subseteq \Gamma$ and $\Delta \cup \{\neg\varphi\}$ is unsatisfiable. Thus, $\Delta \models \varphi$. \square

It is useful to note that the above argument works even when $\neg\varphi \notin \Delta'$.

1.4.1 Compactness using König's Lemma

We present a simple and elegant proof of the compactness theorem that uses König's Lemma. This proof approach works only for propositional logic, and does not extend to first order logic. Let us begin by recalling König's lemma for binary trees.

A binary tree is said to have paths of *arbitrary length* if for each natural number n , there is a path in the tree whose length is $\geq n$. An infinite path in the binary tree is an infinite sequence of vertices of the tree such that successive vertices in the sequence are connected by an edge. Observe that if a binary tree has an infinite path then it also has paths of arbitrary length. This is because for every n , the prefix of the infinite path with $n+1$ vertices, is a path in the binary tree of length n . König's Lemma says that the converse of this is also true.

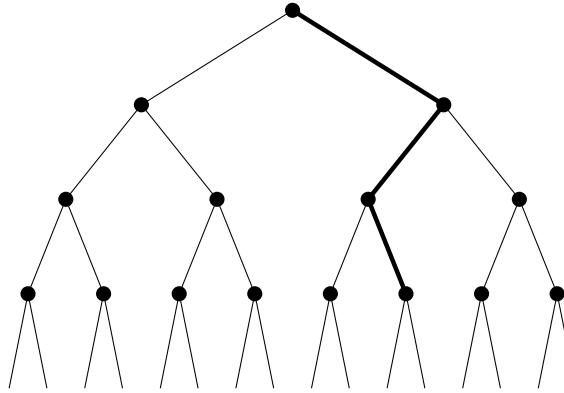


Fig. 1.2 The bold path in the tree, corresponds to the (partial) assignment $v(p_0) = T, v(p_1) = F, v(p_2) = T$.

Lemma 1.28 (König)

A binary tree with paths of arbitrary length has an infinite path.

Proof Suppose u is a vertex that does not have paths of arbitrary length starting from it, then by definition, there must be a number m such that all paths starting from u are of length at most m . Now, if a vertex u has the property that none of its children v have paths of arbitrary length starting from them, then u also cannot have paths of arbitrary length starting from it. The contrapositive of this statement is that if u is vertex with paths of arbitrary length starting from it, then at least one of its children v also has paths of arbitrary length.

Suppose a binary tree has paths of arbitrary length. Then the root is a vertex that has paths of arbitrary length starting from it. The infinite path is given by v_0, v_1, \dots where v_0 is the root. v_{i+1} is the left child of v_i , if the left child has paths of arbitrary length, and v_{i+1} is the right child of v_i otherwise. \square

Let us fix the set of propositions in our logic to be $\text{Prop} = \{p_i \mid i \in \mathbb{N}\}$. Truth assignments to Prop can be thought of as (infinite) paths in the complete, infinite binary tree — vertices at level i correspond to proposition p_i and if the path takes the left child at level i , then it corresponds to the assignment setting p_i to 0; otherwise it sets p_i to 1. Finite paths in this tree, correspond to partial assignments. So a path of length i corresponds to an assignment that sets values to propositions $\{p_0, p_1, \dots, p_{i-1}\}$. For example, in Fig. 1.2, the bold path corresponds to the (partial) assignment v that sets $v(p_0) = T, v(p_1) = F$, and $v(p_2) = T$. Recall that, whether a formula φ holds in an assignment, depends only on the truth values assigned to the propositions that appear in φ (Lemma 1.13). Thus, partial assignments can determine the truth of formulas that only mention the propositions that have been assigned. For example, the partial assignment indicated by the bold path in Fig. 1.2 can determine the truth of any formula that only mentions p_0, p_1 , and p_2 .

Proof (Of Theorem 1.26) Let Γ be a set of formulas that is finitely satisfiable. Logical equivalence (\equiv) partitions Γ into equivalence classes. Take Γ' to be a subset of Γ that contains exactly one representative from each equivalence class. That is, $\Gamma' \subseteq \Gamma$ such that

- For every $\varphi \neq \psi \in \Gamma'$, $\varphi \not\equiv \psi$, and
- For every $\varphi \in \Gamma$, there is $\psi \in \Gamma'$ such that $\varphi \equiv \psi$.

Since $\Gamma' \subseteq \Gamma$, Γ' is also finitely satisfiable.

Recall that $\text{occ}(\varphi)$ is defined to be the set of propositions that appear in φ . For $i \geq 0$, define Γ_i to be

$$\Gamma_i = \{\varphi \in \Gamma' \mid \text{occ}(\varphi) \subseteq \{p_0, p_1, \dots, p_{i-1}\}\}.$$

Observe that Γ_i defines a non-decreasing sequence of sets, i.e., for every i , $\Gamma_i \subseteq \Gamma_{i+1}$. Also, $\Gamma' = \bigcup_{i \geq 0} \Gamma_i$. The most important observation about Γ_i is that it is a finite set — since Γ' has only one formula from each equivalence class of \equiv , each formula in Γ_i corresponds to a unique subset of assignments of $\{p_0, \dots, p_{i-1}\}$ to $\{\text{F}, \text{T}\}$. Thus, we have $|\Gamma_i| \leq 2^{2^i}$. Since Γ' is finitely satisfiable, Γ_i is satisfiable for every i .

Define T_Γ as the following set of (partial) truth assignments.

$$T_\Gamma = \bigcup_{i \geq 0} \{v : \{p_0, \dots, p_{i-1}\} \rightarrow \{\text{F}, \text{T}\} \mid v \models \Gamma_i\}.$$

Recall that we say $v \models \Gamma_i$ if v satisfies all formulas in Γ_i . Consider an assignment $v \in T_\Gamma$ with domain $\{p_0, p_1, \dots, p_{i-1}\}$. By definition $v \models \Gamma_i$. For $j < i$, since $\Gamma_j \subseteq \Gamma_i$, $v \models \Gamma_j$. Further, Γ_j only has propositions $\{p_0, \dots, p_{j-1}\}$, we also have $v' \models \Gamma_j$, where v' is the restriction of v to the domain $\{p_0, \dots, p_{j-1}\}$. So $v' \in T_\Gamma$. Viewed as paths in the infinite binary tree (see Fig. 1.2), v is a path of length i , v' its prefix of length j . What we observe is that T_Γ is “closed” under prefix of paths. Thus, if we restrict our attention to the assignments in T_Γ then they form a subtree of the infinite binary tree.

Let us consider T_Γ . In the previous paragraph we observed that it forms a subtree of the infinite binary tree. It has paths of arbitrary length; this is because every Γ_i is satisfiable, and an (partial) assignment satisfying Γ_i is a path of length i in the tree. Since T_Γ is a binary tree, by König’s lemma, it has an infinite path. The infinite path corresponds to a (full) truth assignment, say v_* . Further, since every prefix of v_* is a (finite) path in T_Γ , it means that the prefix of length i (viewed as a partial assignment) satisfies Γ_i . Therefore, for every i , $v_* \models \Gamma_i$, and hence $v_* \models \Gamma'$. Now, since every formula $\varphi \in \Gamma$ is logically equivalent to some formula $\psi \in \Gamma'$, it means $v_* \models \Gamma$. Thus, Γ is satisfiable. \square

1.4.2 Compactness using Henkin Models

The proof of the compactness theorem we present in the section, relies on constructing a truth assignment for a set of formulas Γ through the process of *saturation*, where we add formulas to the set Γ as long as it remains finitely satisfiable. This is an approach proposed by Henkin.

Let us fix Γ to be a finitely satisfiable set of formulas. We begin by making an important observation about such sets, namely, it can always be extended by adding a formula or its negation, while preserving the property of finite satisfiability.

Lemma 1.29 *Let Γ be finitely satisfiable, and let φ be any formula. Then either $\Gamma \cup \{\varphi\}$ or $\Gamma \cup \{\neg\varphi\}$ is finitely satisfiable.*

Proof Assume (for contradiction), neither $\Gamma \cup \{\varphi\}$ nor $\Gamma \cup \{\neg\varphi\}$ is finitely satisfiable. By definition of finite satisfiability, this means that there are finite subsets $\Gamma_0 \subseteq \Gamma \cup \{\varphi\}$ and $\Gamma_1 \subseteq \Gamma \cup \{\neg\varphi\}$ that are *not* satisfiable. Consider the (finite) set $\Delta = (\Gamma_0 \cup \Gamma_1) \setminus \{\varphi, \neg\varphi\}$. Observe that since $\Delta \subseteq \Gamma$, Δ is satisfiable. Let v be a satisfying truth assignment for Δ . Then either $v \models \varphi$ or $v \models \neg\varphi$. Therefore, either $v \models \Gamma_0$ or $v \models \Gamma_1$, which contradicts our assumption that both Γ_0 and Γ_1 are unsatisfiable. \square

The set of all formulas of propositional logic are *countable*, i.e., there is a 1-to-1, onto function $f : \mathbb{N} \rightarrow \mathcal{F}$, where \mathcal{F} is the set of all propositional logic formulas. Therefore, we can *enumerate* all the formulas in propositional logic. Let $\varphi_0, \varphi_1, \dots$ be an enumeration of all formulas. Let us, inductively, define a sequence of sets of formulas as follows.

$$\begin{aligned} \Delta_0 &= \Gamma \\ \Delta_n &= \begin{cases} \Delta_{n-1} \cup \{\varphi_{n-1}\} & \text{if this is finitely satisfiable} \\ \Delta_{n-1} \cup \{\neg\varphi_{n-1}\} & \text{otherwise} \end{cases} \end{aligned}$$

Observe that the sequence is non-decreasing, i.e., for every n , $\Delta_n \subseteq \Delta_{n+1}$. Further, by induction on n , using Lemma 1.29, we can conclude that Δ_n is finitely satisfiable for all n . Finally, define

$$\Delta = \bigcup_{n \in \mathbb{N}} \Delta_n.$$

Since Δ_n is finitely satisfiable for all $n \in \mathbb{N}$, we can conclude that Δ is also finitely satisfiable.

Proposition 1.30 Δ is finitely satisfiable.

Proof Consider any finite subset $X = \{\psi_1, \dots, \psi_m\}$ of Δ . Observe, by definition Δ , for each i , there is some n_i such that $\psi_i \in \Delta_{n_i}$. Taking $n = \max\{n_1, \dots, n_m\}$, observe that $X \subseteq \Delta_n$. Since Δ_n is finitely satisfiable, X is satisfiable. This means that Δ is finitely satisfiable. \square

Finite satisfiability of Δ implies that Δ is a *complete* set.

Proposition 1.31 *For any formula φ , $\neg\varphi \in \Delta$ if and only if $\varphi \notin \Delta$.*

Proof Without loss of generality assume that φ is the n th formula, i.e., $\varphi = \varphi_n$. Now by definition, in step n of the construction of Δ , if $\varphi \notin \Delta$ then $\neg\varphi \in \Delta$. On the other hand, if $\{\varphi, \neg\varphi\} \subseteq \Delta$ then since $\{\varphi, \neg\varphi\}$ is not satisfiable, Δ would not be finitely satisfiable. But since Δ is finitely satisfiable, it must be the case that at most one out of φ and $\neg\varphi$ belong to Δ . \square

We are now ready to complete the proof of Theorem 1.26. That is, we will show that Γ (which is finitely satisfiable) is satisfiable. Consider the truth assignment v defined as follows.

$$v(p) = \begin{cases} \top & \text{if } p \in \Delta \\ \perp & \text{if } \neg p \in \Delta \end{cases}$$

Note that v is well-defined because by Proposition 1.31, for any proposition p , exactly one among p and $\neg p$ is in Δ . v shows that Δ is satisfiable, because of the following result.

Proposition 1.32 *For any formula $\varphi \in \Delta$, $v \models \varphi$.*

Proof Consider an arbitrary $\varphi \in \Delta$. Let $P = \text{occ}(\varphi)$ and $P^\perp = \{\neg p \mid p \in P\}$. Consider the set

$$U = (\Delta \cap P) \cup (\Delta \cap P^\perp) \cup \{\varphi\}.$$

Since U is a finite subset of Δ , by Proposition 1.30, we have U is satisfiable. Let v' be a truth assignment such that $v' \models U$. Observe that for every $p \in \Delta \cap P$, $v'(p) = \top$ and for every $\neg p \in \Delta \cap P^\perp$, $v'(\neg p) = \perp$. Therefore, v and v' agree on all propositions in P . By Lemma 1.13, since $v' \models \varphi$, we have $v \models \varphi$. \square

Proposition 1.32 establishes the fact that $v \models \Delta$. Since $\Gamma \subseteq \Delta$, $v \models \Gamma$. Therefore, Γ is satisfiable.

1.4.3 An Application of Compactness: Coloring Infinite Planar Graphs

In this section we present an application of the compactness theorem. We will show that all infinite planar graphs are 4-colorable. We begin by recalling the graph coloring problem, and its connection to propositional logic.

Definition 1.33 (Graphs)

An undirected graph $G = (V, E)$ is a set of vertices V , and a set of edges $E \subseteq V \times V$, such that E is symmetric (i.e., $(u, v) \in E$ iff $(v, u) \in E$) and irreflexive (i.e., $(u, u) \notin E$ for any $u \in V$).

Definition 1.34 (Coloring)

A k -coloring of graph $G = (V, E)$ is a function $c : V \rightarrow \{1, 2, \dots, k\}$ such that if $(u, v) \in E$ then $c(u) \neq c(v)$. If G has a k -coloring then G is said to be k -colorable.

The problem of determining whether a graph is k -colorable can be “reduced” to checking the satisfiability of a set of formulas.

Proposition 1.35 For any graph $G = (V, E)$ (with possibly infinitely many vertices), there is a set of formulas $\Gamma_{G,k}$ such that G is k -colorable iff $\Gamma_{G,k}$ is satisfiable.

Proof For each vertex $u \in V$ and $1 \leq i \leq k$, take the proposition r_{ui} to denote “vertex u has color i ”. $\Gamma_{G,k}$ is the following set of formulas.

- For each $u \in V$, the formula $r_{u1} \vee r_{u2} \vee \dots \vee r_{uk}$. Intuitively these formulas capture the constraint that every vertex gets at least one of the k colors.
- For each $u \in V$ and $1 \leq i, j \leq k$ with $i \neq j$, the formula $\neg r_{ui} \vee \neg r_{uj}$. These formulas capture the constraint that a vertex does not get two different colors.
- For each edge $(u, v) \in E$ and color $1 \leq i \leq k$, the formula $\neg r_{ui} \vee \neg r_{vi}$. These formulas ensure that adjacent vertices do not get the same color.

For a coloring c , define the valuation v_c such that $v_c(r_{ui}) = \top$ iff $c(u) = i$. Similarly for a valuation v , define a function $c_v(u) = i$ iff $v(r_{ui}) = \top$. Observe that

- If c is a valid k -coloring of G then v_c satisfies $\Gamma_{G,k}$, and
- If v satisfies $\Gamma_{G,k}$ then c_v is a valid k -coloring of G .

Proof of the above observations is left as exercise. \square

Finite, planar graphs are graphs with finitely many vertices such that there is a drawing of the graph on the plane where the edges do not cross. A celebrated result about finite, planar graphs is that 4 colors are sufficient to color the graph.

Theorem 1.36 (Appel-Haken)

Every finite planar graph is 4-colorable.

We will show that the compactness theorem in fact shows that Theorem 1.36 can be extended to infinite graphs as well.

Corollary 1.37 *All infinite planar graphs are 4-colorable.*

Proof Let G be an infinite planar graph. Consider the set of formulas $\Gamma_{G,4}$ constructed in Proposition 1.35. Observe that $\Gamma_{G,4}$ is finitely satisfiable. This can be seen as follows. Let Γ_0 be any finite subset of $\Gamma_{G,4}$. Let G_0 be the graph induced by the vertices u such that the proposition r_{ui} appears in Γ_0 for some i . Now, G_0 is a finite, planar graph and so by Theorem 1.36 has a 4-coloring c . Then by the proof of Proposition 1.35, the valuation v_c satisfies $\Gamma_{G_0,4}$. Since $\Gamma_0 \subseteq \Gamma_{G_0,4}$, we have v_c satisfies Γ_0 .

Since $\Gamma_{G,4}$ is finitely satisfiable, by the compactness theorem, $\Gamma_{G,4}$ is satisfiable. Let v be a satisfying assignment for $\Gamma_{G,4}$. Again, by Proposition 1.35, c_v is a valid 4-coloring of G . \square

Chapter 2

Proof Systems

If logic is the science of valid inference, then proofs embody its heart. But what are mathematical proofs? They are a sequence of statements where each statement in the sequence is either a self evident truth, or “logically” follows from previous observations. Thus, sound derivation principles are identified by correct proofs.

Example 2.1 Euclid’s Elements sets out axioms (or postulates), which are self evident truths, and proves all results in geometry from these truths formally. Euclid lays out five axioms for geometry.

- A1 A straight line can be drawn from any point to any point.
- A2 A finite line segment can be extended to an infinite straight line.
- A3 A circle can be drawn with any point as center and any given radius.
- A4 All right angles are equal.
- A5 If a straight line falling on two straight lines makes the interior angles on the same side less than two right angles, the straight lines, if produced indefinitely, will meet on that side on which the angles are less than two right angles.

Using these axioms, Euclid proves a number of results in geometry. He uses previously proved propositions in the proofs of later observations. An example of such a result, is the proof that the sum of the interior angles of a triangle is 180° .

Claim The interior angles of a triangle sum to two right angles.

Proof Consider the diagram in Fig. 2.1. The proposition is proved using the following sequence of statements.

1. Extend one side (say) BC to D [A2]
2. Draw a line parallel to AB through point C; call it CE [P31]
3. Since AB is parallel to CE, $BAC = ACE$ and $ABC = ECD$ [P29]
4. Thus, the sum of the interior angles = $ACB + ACE + ECD = 180^\circ$

References [P31] and [P29] in steps 2 and 3, allude to previously propositions 31 and 29, proved in the book. \square

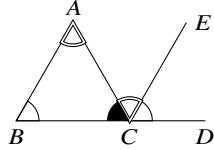


Fig. 2.1 Proof that the sum of the internal angles of a triangle are 180°

Example 2.1 highlights the basic elements of identifying good proofs — one needs to identify axioms, and the principle by which new conclusions can be drawn from previously established facts. A formal proof system for a logic identifies such *axioms* and *rules of inference*. We will introduce two such proof systems for propositional logic — a Frege-style proof system, and resolution — to give a flavor of different types of proof systems.

2.1 A Frege-style Proof System

Proof systems are most convenient presented as a collection of rules of the form

$$\frac{\Gamma}{\varphi}$$

where Γ is a set of formulas (schemas) and φ is a formula (schema). Such rules can be interpreted as follows — if every formula in Γ can be established then φ can be concluded from these observations. One special case is when $\Gamma = \emptyset$. In this case the formula below the line can be concluded without establishing anything; in other words, it is an axiom. Instead of explicitly writing \emptyset above the line, we simply don't write anything, and present this axiom in the form

$$\overline{\varphi}$$

Before presenting our first proof system, observe that all propositional logic formulas can be expressed using just implication and \perp . To see this observe that $\neg\varphi$ is the same as $\varphi \rightarrow \perp$ and $\varphi \vee \psi$ is $(\varphi \rightarrow \perp) \rightarrow \psi$. Our first proof system, shown in Fig. 2.2, assumes that all formulas are written using implication and \perp .

Our first proof system has 3 axiom *schemas* and one rule of inference. The formulas φ , ψ , and ρ in Fig. 2.2, can be *any formulas*. For example, taking $\varphi = p$ and $\psi = p$, we get $p \rightarrow (p \rightarrow p)$ as an *instantiation* of the first axiom schema, while taking $\varphi = p$ and $\psi = p \rightarrow p$, we get $p \rightarrow ((p \rightarrow p) \rightarrow p)$ as a different instantiation of the *same* schema. The rule of inference in this proof system, is a very commonly used rule. It, therefore, has a special name; it is called *modus ponens*.

$$\begin{array}{c}
 \frac{}{\varphi \rightarrow (\psi \rightarrow \varphi)} \quad \frac{}{(\varphi \rightarrow (\psi \rightarrow \rho)) \rightarrow ((\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \rho))} \\
 \frac{}{((\varphi \rightarrow \perp) \rightarrow \perp) \rightarrow \varphi} \quad \frac{\varphi \quad \varphi \rightarrow \psi}{\psi}
 \end{array}$$

Fig. 2.2 A Frege-style Proof System

Proofs in our (formal) proof system, will be like the usual proofs in mathematics — they will be a sequence of statements. However, instead of using english statements, here they will simply be well formed formulas of propositional logic. The statement (or formula) being proved is the last one in the sequence. The sequence of formulas in a proof should be consistent with the axioms and rule of inference of the proof system, for it to be *valid* proof. This is captured in the definition below.

Definition 2.2 (Proofs)

A *proof* of φ from a set (possibly infinite) of hypotheses Γ is a finite sequence of wffs $\psi_1, \psi_2, \dots, \psi_m$ such that $\psi_m = \varphi$, and for every $k \in \{1, 2, \dots, m\}$, either

- $\psi_k \in \Gamma$, or
- ψ_k is an axiom, or
- ψ_k follows from ψ_i and ψ_j , with $i, j < k$, by modus ponens.

The *length* of such a proof is the number of wffs in the sequence, namely, m . If there is a proof of φ from Γ , we denote this by $\Gamma \vdash \varphi$. When $\Gamma = \emptyset$, we write this as $\vdash \varphi$ (as opposed to $\emptyset \vdash \varphi$).

Let us look at some proofs in our system.

Example 2.3 Let us construct a proof of $q \rightarrow p$ from the hypothesis $\{p\}$. Such a proof is as follows.

1. $p \rightarrow (q \rightarrow p)$ Axiom 1, taking $\varphi = p$, and $\psi = q$
2. p Hypothesis in set Γ
3. $q \rightarrow p$ Modus Ponens on lines 1 and 2

Thus, $\{p\} \vdash q \rightarrow p$.

We will now show that $\vdash \perp \rightarrow ((p \rightarrow q) \rightarrow (p \rightarrow p))$.

1. $(p \rightarrow (q \rightarrow p)) \rightarrow ((p \rightarrow q) \rightarrow (p \rightarrow p))$
Axiom 2, taking $\varphi = p$, $\psi = q$ and $\rho = p$
2. $p \rightarrow (q \rightarrow p)$
Axiom 1, taking $\varphi = p$, and $\psi = q$
3. $(p \rightarrow q) \rightarrow (p \rightarrow p)$
Modus ponens on lines 1 and 2
4. $((p \rightarrow q) \rightarrow (p \rightarrow p)) \rightarrow (\perp \rightarrow ((p \rightarrow q) \rightarrow (p \rightarrow p)))$
Axiom 1, taking $\varphi = (p \rightarrow q) \rightarrow (p \rightarrow p)$, and $\psi = \perp$
5. $\perp \rightarrow ((p \rightarrow q) \rightarrow (p \rightarrow p))$
Modus ponens on lines 3 and 4

Finally, let us show $\vdash p \rightarrow p$.

1. $(p \rightarrow ((p \rightarrow p) \rightarrow p)) \rightarrow ((p \rightarrow (p \rightarrow p)) \rightarrow (p \rightarrow p))$
Axiom of 2, taking $\varphi = p$, $\psi = p \rightarrow p$ and $\rho = p$
2. $(p \rightarrow ((p \rightarrow p) \rightarrow p))$
Axiom 1, taking $\varphi = p$, and $\psi = p \rightarrow p$
3. $(p \rightarrow (p \rightarrow p)) \rightarrow (p \rightarrow p)$
Modus ponens on lines 1 and 2
4. $p \rightarrow (p \rightarrow p)$
Axiom 1, taking $\varphi = p$, $\psi = p$
5. $p \rightarrow p$
Modus ponens on lines 3 and 4

In proof systems, like the one we are considering in this section, there is a very useful theorem that makes writing proofs easy. This is called the *deduction theorem*. Some proof systems have it as an explicit rule.

Theorem 2.4 (Deduction Theorem)

If $\Gamma \cup \{\varphi\} \vdash \psi$ then $\Gamma \vdash \varphi \rightarrow \psi$.

First, observe that the converse of Theorem 2.4, is clearly true, i.e., if $\Gamma \vdash \varphi \rightarrow \psi$ then $\Gamma \cup \{\varphi\} \vdash \psi$. Establishing this left as an exercise. The proof of the deduction theorem is a more difficult exercise. The informal outline of the proof is as follows. Assume that $\rho_1, \rho_2, \dots, \rho_m$ is a proof of ψ from $\Gamma \cup \{\varphi\}$. One shows by induction on i that, for each line i , we have $\Gamma \vdash \varphi \rightarrow \rho_i$.

The deduction theorem simplifies the task of writing down proofs in our proof system.

Example 2.5 Consider the task of showing $\vdash (\varphi \rightarrow \psi) \rightarrow ((\psi \rightarrow \rho) \rightarrow (\varphi \rightarrow \rho))$, where φ , ψ , and ρ are arbitrary wffs. Our approach to solving this problem, would instead be to instead establish $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \rho$. If we succeed, we will get the desired result by using the deduction theorem a few times.

1. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \varphi \rightarrow \psi$ Hypothesis
2. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \varphi$ Hypothesis
3. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \psi$ Modus Ponens on lines 1 and 2
4. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \psi \rightarrow \rho$ Hypothesis
5. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho, \varphi\} \vdash \rho$ Modus Ponens on lines 3 and 4
6. $\{\varphi \rightarrow \psi, \psi \rightarrow \rho\} \vdash \varphi \rightarrow \rho$ Deduction Theorem
7. $\{\varphi \rightarrow \psi\} \vdash (\psi \rightarrow \rho) \rightarrow (\varphi \rightarrow \rho)$ Deduction Theorem
8. $\vdash (\varphi \rightarrow \psi) \rightarrow ((\psi \rightarrow \rho) \rightarrow (\varphi \rightarrow \rho))$ Deduction Theorem

2.1.1 Completeness Theorem

Proofs in a formal proof system like the one in Fig. 2.2 try capture the notion of correct logical inference. But what do we mean by “correct logical inference”? There

are two aspects to such a question. First, any conclusion drawn by a proof must be logically correct, i.e., consistent with the semantics we defined. In other words, we should not be able to conclude any false facts using a proof. This is often referred to as the *soundness* of the proof system. The second is that the proof system must be rich enough to be able to prove all true facts. This is called the *completeness* of the proof system. We make this connection between provability and the semantics we gave precise in the following theorem.

Theorem 2.6 (Soundness and Completeness)

For any set of formulas Γ (possibly even infinite) and any wff φ the following two properties hold.

Soundness *If $\Gamma \vdash \varphi$ then $\Gamma \models \varphi$.*

Completeness *If $\Gamma \models \varphi$ then $\Gamma \vdash \varphi$.*

Thus, any formula proved without any hypotheses is a tautology, and every tautology has a proof from the empty set of hypotheses in our proof system. We will not prove Theorem 2.6. We will instead prove such a soundness and completeness theorem for the second proof system that we will introduce for propositional logic. Proving soundness is usually easy. It requires making sure that the axioms and proof rules are consistent with the semantics of the logic. In this case it requires showing that every axiom in the proof system is indeed a tautology, and modus ponens is consistent with logical consequence. Proving completeness is typically hard.

2.2 Resolution

Notice that the proofs for formulas that we constructed in Examples 2.3 and 2.5 are very symbolic and mechanical — one doesn't need to understand what the formula we are trying to prove is saying or what the meaning of the hypotheses is. Instead the proofs are constructed by looking at the pattern of formulas. This raises the prospect of trying to mechanize the process of searching for proofs of formulas. However, the proof system in Sect. 2.1 is not good for this purpose. This is because at any point during the construction of the proof, one can extend it by using any one of the axiom schemas. Since each axiom schema can be instantiated in infinitely many possible ways, this makes mechanization difficult. A proof system that is amenable to mechanization is one that has very few choices at any step of the proof construction. *Resolution* is such a proof system. Resolution has no axioms and only one rule of inference.

Resolution is a method for *refutations*, i.e., it proves that a formula is “not true” or more precisely *unsatisfiable*; recall that a formula φ is unsatisfiable if $\mathbf{v} \not\models \varphi$ for all valuations \mathbf{v} . Thus unlike the proof system in Sect. 2.1 which proves validity of a formula directly, resolution works by showing that the negation of the formula is unsatisfiable. One can imagine resolution refutations like a proof by contradiction, where one assumes the negation of what one is trying to establish, and trying to show that that is impossible.

Resolution works when the formulas are represented in *conjunctive normal form* (CNF). We begin, therefore, by introducing conjunctive normal form formulas. CNF formulas are built using propositions and their negations, and the logical operators of \wedge and \vee . It will be convenient to think of \vee and \wedge being present implicitly (see Definition 2.7 below) instead of explicitly in the syntax.

Definition 2.7 (Conjunctive Normal Form)

Conjunctive normal form formulas are defined as follows.

- A *literal* is a proposition p or its negation $\neg p$.
- A *clause* is a disjunction of literals. We will think of a clause as a set of literals, implicitly assuming that the literals are disjuncted. In this interpretation, a truth valuation satisfies a clause if some literal in the set evaluates to 1 under the truth valuation.
- The *empty clause* is the clause containing no literals. By definition, no truth assignment satisfies the empty clause.
- A formula is said to be in *conjunctive normal form* (CNF) if it is conjunction of clauses. Again, we will think of a CNF formula as a set of clauses, with the conjunction being implicit in the syntax. With this interpretation, a truth valuation satisfies a CNF formula if it satisfies every clause in the set representing the formula.

CNF formulas as sets of sets of literals

In Definition 2.7, clauses are represented as sets of literals, and CNF formulas as sets of clauses. Why is this well-defined? The reason is because \vee and \wedge are both *idempotent*, *commutative*, and *associative*. Idempotence means that disjuncting/conjuncting a formula with itself is equivalent to the formula. In other words, for any wff φ we have

$$\varphi \vee \varphi \equiv \varphi \quad \varphi \wedge \varphi \equiv \varphi.$$

Idempotence of disjunction and conjunction ensures that representing them as sets which don't have repeated elements is faithful with the semantics. Commutativity and associativity mean that the order in which formulas are disjuncted/conjuncted does not change its semantics. That is, for any formulas φ , ψ , and ρ ,

$$\begin{aligned} \varphi \vee \psi &\equiv \psi \vee \varphi & \varphi \wedge \psi &\equiv \psi \wedge \varphi \\ (\varphi \vee \psi) \vee \rho &\equiv \varphi \vee (\psi \vee \rho) & (\varphi \wedge \psi) \wedge \rho &\equiv \varphi \wedge (\psi \wedge \rho). \end{aligned}$$

Thus, commutativity and associativity ensure that sets (which are ordered collections) are faithful representations of disjunctions and conjunctions of formulas.

CNF formulas are formulas in a restricted form. However, they are not semantically restrictive. That is, every wff φ can be shown to be equivalent to a formula ψ in CNF — we can push negations all the way in using DeMorgan's Laws, and then

distribute \vee over \wedge . We will look at this conversion process more closely later in this chapter. Let us look at some examples of CNF and non-CNF formulas.

Example 2.8 The formulas $(p_1 \wedge q_1) \vee (p_2 \wedge q_2)$, $\neg(p \wedge q)$ are examples of formulas that are not in CNF — neither formula has \wedge as the topmost connective. The formulas $(p_1 \vee p_2) \wedge (p_1 \vee q_2) \wedge (q_1 \vee p_2) \wedge (q_1 \vee q_2)$ and $\neg p \vee \neg q$ are formulas in CNF as they are conjunctions of clauses.

We will represent CNF formulas as set of set of literals, without explicit conjunctions and disjunctions. For example, $\{\{p_1, p_2\}, \{p_1, q_2\}, \{q_1, p_2\}, \{q_1, q_2\}\}$ is the way the formula $(p_1 \vee p_2) \wedge (p_1 \vee q_2) \wedge (q_1 \vee p_2) \wedge (q_1 \vee q_2)$ will be represented. Similarly, $\{\{\neg p, \neg q\}\}$ will be the representation of $\neg p \vee \neg q$.

The resolution proof system is a sequence of transformations that preserve satisfiability, until the empty clause (which is by definition not satisfiable) is obtained. The transformations involve the single rule of inference that constructs the *resolvent* of two clauses.

Definition 2.9 (Resolvent) The only rule in the resolution proof system is as follows.

$$\frac{C \cup \{p\} \quad D \cup \{\neg p\}}{C \cup D}$$

The conclusion $C \cup D$ is called the *resolvent* of $C \cup \{p\}$ and $D \cup \{\neg p\}$ with respect to proposition p .

Two clauses may have multiple resolvents depending on which proposition one chooses to resolve with respect to. The resolvent of two clauses may be the empty clause if C and D are empty sets. Let us look at some examples to clarify this definition.

Example 2.10 Consider the clauses $\{p, \neg q, \neg r\}$ and the clause $\{\neg p, \neg q\}$. The resolvent of these two clauses (with respect to p) is the clause $\{\neg q, \neg r\}$.

On the other hand, if we consider clauses $\{p, \neg q\}$ and $\{\neg p, q\}$, we have two possible resolvents. If we resolve with respect to p , we get $\{q, \neg q\}$, and if we resolve with respect to q , we get $\{p, \neg p\}$.

Definition 2.11 (Refutations)

A *resolution refutation* of a (possibly infinite) set of clauses Γ is a sequence of clauses C_1, C_2, \dots, C_m such that each clause C_k is either in Γ or a resolvent of two clauses C_i and C_j ($i, j < k$), and the last clause C_m in the refutation is the empty clause.

Example 2.12 The set of clauses $\Gamma = \{\{p, q\}, \{\neg p, r\}, \{\neg q, r\}, \{\neg r\}\}$ has the following resolution refutation.

1. $\{\neg p, r\}$
2. $\{\neg r\}$
3. $\{\neg p\}$ Resolvent of 1 and 2
4. $\{\neg q, r\}$
5. $\{\neg q\}$ Resolvent of 2 and 4
6. $\{p, q\}$
7. $\{q\}$ Resolvent of 3 and 6
8. $\{\}$ Resolvent of 5,7

2.2.1 Proving Tautologies with Resolution

As mentioned before, resolution refutations establish the unsatisfiability of a formula given in CNF. To prove that a formula is a tautology using resolution, we need to use Proposition 1.24. That is, to prove that φ is a tautology, we need to convert $\neg\varphi$ into CNF. This can be done as follows.

1. Push negations inside using DeMorgan's Laws. Recall that DeMorgan's laws say the following.

$$\neg(\psi_1 \wedge \psi_2) \equiv \neg\psi_1 \vee \neg\psi_2 \quad \neg(\psi_1 \vee \psi_2) \equiv \neg\psi_1 \wedge \neg\psi_2$$

2. Remove double negations, because $\neg\neg\psi \equiv \psi$
3. Distribute \vee over \wedge , using the distributive law, which says

$$\psi_1 \vee (\psi_2 \wedge \psi_3) \equiv (\psi_1 \vee \psi_2) \wedge (\psi_1 \vee \psi_3) \quad (\psi_1 \wedge \psi_2) \vee \psi_3 \equiv (\psi_1 \vee \psi_3) \wedge (\psi_2 \vee \psi_3)$$

Example 2.13 Let $\varphi = (\neg p_1 \vee \neg q_1) \wedge (\neg p_2 \vee \neg q_2)$. We can convert $\neg\varphi$ to CNF as follows.

1. Pushing negations inside using DeMorgan's Laws, we get

$$(\neg\neg p_1 \wedge \neg\neg q_1) \vee (\neg\neg p_2 \wedge \neg\neg q_2)$$

2. Removing double negations, we get

$$(p_1 \wedge q_1) \vee (p_2 \wedge q_2)$$

3. Distributing \vee over \wedge , we get

$$(p_1 \vee p_2) \wedge (p_1 \vee q_2) \wedge (q_1 \vee p_2) \wedge (q_1 \vee q_2)$$

The above method while semantically correct, can be expensive. The main reason is that when we distribute \vee over \wedge , the resulting formula can be exponentially larger. For example, if we have the formula

$$\bigvee_{i=1}^n (p_i \wedge q_i)$$

and we distribute \vee over \wedge , we will get a CNF formula where each clause is of the form $(r_1 \vee r_2 \vee \dots \vee r_n)$, where r_i is either p_i or q_i . This will result in a formula with 2^n clauses (as we have two choices for each r_i).

The exponential blowup can be avoided by using a translation proposed by Tseitin. Tseitin's method does not construct a logically equivalent CNF formula. Instead, for the formula $\neg\varphi$, it constructs a CNF formula ψ with the property that $\neg\varphi$ is satisfiable if and only if ψ is satisfiable. This weaker correspondence between $\neg\varphi$ and ψ is sufficient in this context; we will have φ is a tautology if and only if ψ is unsatisfiable.

Tseitin's Method.

Let us describe the conversion of formula $\neg\varphi$. The first step is to introduce new *extension* propositions x_ψ for each subformula ψ of φ as follows.

- For each proposition in φ , the extension proposition x_p is the same as p .
- For each negated subformula $\neg\psi$, $x_{\neg\psi}$ is taken to be the literal $\neg x_\psi$.
- For all other subformulas ψ , x_ψ is a new proposition.

Having identified the extension propositions, the CNF formula that we will construct corresponding to $\neg\varphi$ is as follows. Here, we will use the representation of CNF formulas as sets of sets of literals. So for φ , we define Γ_φ to be the following set of clauses.

- The singleton clause $\{\neg x_\varphi\}$
- For each subformula $\psi \wedge \rho$, we add the CNF formula equivalent to " $x_{\psi \wedge \rho} \leftrightarrow x_\psi \wedge x_\rho$ ". In other words, we will have the clauses

$$\{\neg x_{\psi \wedge \rho}, x_\psi\} \quad \{\neg x_{\psi \wedge \rho}, x_\rho\} \quad \{x_{\psi \wedge \rho}, \neg x_\psi, \neg x_\rho\}$$

- For each subformula $\psi \vee \rho$, we add the CNF formula equivalent to " $x_{\psi \vee \rho} \leftrightarrow x_\psi \vee x_\rho$ ". In other words, we will have the clauses

$$\{x_{\psi \vee \rho}, \neg x_\psi\} \quad \{x_{\psi \vee \rho}, \neg x_\rho\} \quad \{\neg x_{\psi \vee \rho}, x_\psi, x_\rho\}$$

Observe that the number of subformulas of a given formula φ is linear in the size of φ . Thus the number of extension propositions we introduce is linear in φ . Further, for each subformula, we are introducing only a constant (3 to be precise) number of clauses. This means that resulting set of clauses Γ_φ is linear in the size of φ .

Example 2.14 Let us apply Tseitin's construction to the formula in Example 2.13. Recall that $\varphi = (\neg p_1 \vee \neg q_1) \wedge (\neg p_2 \vee \neg q_2)$. The first step is to identify the new extension propositions we need. In this case there are only 3 that we will add — x_φ corresponding to φ , x_1 corresponding to $(\neg p_1 \vee \neg q_1)$, and x_2 corresponding to

$(\neg p_2 \vee \neg q_2)$. Our CNF formula Γ_φ will be obtained by adding 3 clauses for each of these interesting subformulas corresponding to x_φ , x_1 and x_2 . Thus, we have Γ_φ is the following set

$$\left\{ \begin{array}{l} \{\neg x_\varphi\}, \\ \{\neg x_\varphi, x_1\}, \{\neg x_\varphi, x_2\}, \{x_\varphi, \neg x_1, \neg x_2\}, \\ \{x_1, p_1\}, \{x_1, q_1\}, \{\neg x_1, \neg p_1, \neg q_1\}, \\ \{x_2, p_2\}, \{x_2, q_2\}, \{\neg x_2, \neg p_2, \neg q_2\} \end{array} \right\}$$

In the above description, we have replaced $\neg\neg p$ by p for $p \in \{p_1, q_1, p_2, q_2\}$.

2.2.2 Completeness of Resolution

We will now prove that resolution is a “correct” proof system. In other words, we will prove the soundness and completeness of resolution. What that means for resolution is that a set of clauses Γ has a resolution refutation if and only if Γ is unsatisfiable. Note that we will establish this for any set Γ , including those that are infinite. Recall that a set of clauses Γ is satisfiable if there is a truth assignment v such that for every clause $C \in \Gamma$, there is some literal $\ell \in C$ such that $v[\ell] = T$.

Theorem 2.15 (Soundness)

If a set of clauses Γ has a resolution refutation, then Γ is unsatisfiable.

Proof The crux of the proof of the soundness theorem, is to establish the “correctness” of the resolution proof rule. That is captured by the following lemma.

Lemma 2.16 (Resolution Lemma)

Let Δ be a set of clauses and let C be the resolvent of two clauses $D, E \in \Delta$. Then for any assignment v , $v \models \Delta$ if and only if $v \models \Delta \cup \{C\}$. \square

Proof (Of Lemma 2.16) Observe that if $v \models \Delta \cup \{C\}$, then clearly $v \models \Delta$.

Consider a truth assignment such that $v \models \Delta$. Without loss of generality, let us assume there is a proposition p such that $p \in D$ and $\neg p \in E$. If $v(p) = T$ then since $v \models E$, there must be a literal $\ell \in E$ (obviously $\ell \neq \neg p$) such that $v[\ell] = T$. Since $\ell \in C$, we have $v \models C$. On the other hand, if $v(p) = F$ then since $v \models D$, there must be a literal $\ell \in D$ (obviously $\ell \neq p$) such that $v[\ell] = T$. Since $\ell \in C$, we have $v \models C$. \square

Using Lemma 2.16, we are ready to complete the proof of Theorem 2.15. Let C_1, C_2, \dots, C_m be a resolution refutation of Γ . Let us define a sequence of sets of clauses inductively as follows.

$$\Gamma_0 = \Gamma \quad \Gamma_i = \Gamma_{i-1} \cup \{C_i\}$$

Thus, $\Gamma_i = \Gamma \cup \{C_1, C_2, \dots, C_i\}$. Since $C_m = \{\}$, Γ_m is clearly unsatisfiable. Therefore, by Lemma 2.16 (and induction), $\Gamma = \Gamma_0$ is also unsatisfiable. \square

Theorem 2.17 (Completeness)

If a set of clauses Γ is unsatisfiable then there is a resolution refutation of Γ .

We will present a proof of the completeness theorem due to David and Putnam. For finite Γ , the proof is constructive. That is, when Γ is unsatisfiable, it gives a specific construction of a refutation for Γ .

Proof Let Γ be an unsatisfiable set of clauses. By the compactness theorem (Theorem 1.26), there is a finite subset $\Delta \subseteq \Gamma$ that is unsatisfiable; this can be seen by taking $\varphi = \perp$ in Corollary 1.27. We will prove the completeness theorem by induction on the number of propositions appearing in Δ . Before outlining the proof, it is useful to point out that since Δ is unsatisfiable, it must be a *non-empty* set of clauses. This is because an empty set of clauses is (by definition) satisfiable.

For the base case, observe that if Δ contains no propositions, then Δ contains the empty clause. Then the refutation for Δ is simply $\{\}$.

Let us now consider the induction step. Let us call a clause C *trivial* if there is a proposition p such that $\{p, \neg p\} \subseteq C$. Trivial clauses are valid, and hence they can be removed from Δ without affecting its satisfiability. Thus, without loss of generality, we will assume that Δ does not have any trivial clauses. Let p be a proposition that appears in Δ . With respect to proposition p , Δ can be partitioned into 3 sets.

$$\begin{aligned}\Delta_0^p &= \{C \in \Delta \mid C \cap \{p, \neg p\} = \emptyset\} \\ \Delta_+^p &= \{C \in \Delta \mid p \in C\} \\ \Delta_-^p &= \{C \in \Delta \mid \neg p \in C\}\end{aligned}$$

Thus, Δ_0^p are clauses where p does not appear, Δ_+^p are those where p appears positively, and Δ_-^p are those where p appears negatively. Let us construct a new set of clauses as follows.

$$\Delta_p = \Delta_0^p \cup \{C \cup D \mid C \cup \{p\} \in \Delta_+^p \text{ and } D \cup \{\neg p\} \in \Delta_-^p\}$$

Thus, Δ_p has all the clauses in Δ_0^p and all resolvents of clauses from Δ_+^p and Δ_-^p . Observe that p no longer appears in Δ_p . If we can argue that Δ_p is unsatisfiable then we can complete the proof by using the induction hypothesis — the refutation for Δ is just all the steps to create Δ_p followed by a refutation for Δ_p .

To finish the proof, we need to establish the following lemma.

Lemma 2.18 If Δ_p is satisfiable then so is Δ . □

Proof (Of Lemma 2.18) Let v be a truth assignment that satisfies Δ_p . Let v' be the truth assignment that is identical to v , except that it flips the assignment to p . Observe that since v and v' only differ on the assignment to p , they agree on all the propositions appearing in Δ_p . Therefore, v' also satisfies Δ_p .

Let us assume without loss of generality, that $v(p) = \top$ and $v'(p) = \perp$. We will show that either v or v' satisfies Δ . Observe that both v and v' satisfy Δ_0^p (because p does not appear in Δ_0^p). Also, v satisfies Δ_+^p (because all clauses in Δ_+^p have p) and v' satisfies Δ_-^p (because all clauses in Δ_-^p have $\neg p$). Now if v satisfies Δ_-^p , v satisfies

Δ . Similarly, if ν' satisfies Δ_+^P then ν' satisfies Δ . So the problem is if neither of these hold. In that case that is a clause $C \cup \{p\} \in \Delta_+^P$ that is not satisfied by ν' and there is a clause $D \cup \{\neg p\} \in \Delta_-^P$ that is not satisfied by ν . But then their resolvent $C \cup D \in \Delta_P$ is not satisfied by either ν or ν' , which contradicts our assumption that both ν and ν' satisfy Δ_P . \square

With the proof of Lemma 2.18, we have completed the proof of Theorem 2.17. \square

2.3 Craig's Interpolation Theorem and Proof Complexity

Craig's Interpolation theorem is a classical result in logic that holds for many different logics. The theorem has been used in different contexts in the broad area of formal methods and verification — in hardware and software specification; reasoning about large knowledge databases; type inference; combination of theorem provers for different first order theories; model checking of finite and infinite state systems through the construction of abstractions. In this section we look at some of its connections to theoretical computer science and complexity theory in particular. We will introduce the theorem for propositional logic, and its connection with proof lengths for propositional logic formulas.

2.3.1 Craig's Interpolation Theorem

Before we state and prove the interpolation theorem, it will be convenient to introduce some notation. A list of propositions p_1, p_2, \dots, p_n will be denoted by \vec{p} when the actual number of propositions in the list is unimportant. A formula φ over propositions $p_1, p_2, \dots, p_n, q_1, q_2, \dots, q_m$ will be sometimes denoted as $\varphi(\vec{p}, \vec{q})$, making explicit the propositions that *may* syntactically appear in φ . Finally, for a truth valuation ν and a list of propositions \vec{p} , $\nu \upharpoonright_{\vec{p}}$ will denote the restriction of ν to the propositions \vec{p} ; note that ν has an infinite domain, the domain of $\nu \upharpoonright_{\vec{p}}$ is finite and restricted to \vec{p} .

Theorem 2.19 (Craig)

Suppose $\models \varphi(\vec{p}, \vec{q}) \rightarrow \psi(\vec{q}, \vec{r})$. Then there is a formula $\eta(\vec{q})$ such that $\models \varphi(\vec{p}, \vec{q}) \rightarrow \eta(\vec{q})$ and $\models \eta(\vec{q}) \rightarrow \psi(\vec{q}, \vec{r})$. η is said to be the interpolant of φ and ψ .

Before presenting the proof of Theorem 2.19, let us examine its statement. There are different equivalent ways of presenting this result. Recall that $\models \varphi \rightarrow \psi$ iff $\varphi \models \psi$. Therefore, we could say that if ψ is a logical consequence of φ then η is a formula that captures the reason why, using only the common propositions. In this case, since $\varphi \models \eta$, we could think of η as an “abstraction” of φ (as η “forgets” the constraints φ imposes on \vec{p}) that is sufficient to ensure ψ . Another formulation of Craig's theorem

is as follows. Suppose $\varphi(\vec{p}, \vec{q}) \wedge \psi(\vec{q}, \vec{r})$ is unsatisfiable (or equivalently, $\varphi \models \neg\psi$), then there is a formula $\eta(\vec{q})$ over the common propositions such that $\varphi \models \eta$ and $\eta \wedge \psi$ is unsatisfiable (i.e., $\eta \models \neg\psi$). Informally, η is an abstraction of φ that captures why $\varphi \wedge \psi$ is unsatisfiable. We will consider this formulation of Craig's theorem in terms of unsatisfiability, when we revisit resolution later in this section.

Proof (Of Theorem 2.19) The proof of Craig's Interpolation Theorem is quite simple in this context of propositional logic. Let us define

$$M = \{v \upharpoonright_{\vec{q}} \mid v[\varphi] = T\}.$$

Observe that since the domain (and range) of all functions in M is finite, M is a finite set. Let us, without loss of generality, take $M = \{v_1, v_2, \dots, v_n\}$. The interpolant η will essentially say that “one among the assignments in M hold”. Formally,

$$\eta(\vec{q}) = \vee_{i=1}^n (q_1^{(i)} \wedge q_2^{(i)} \wedge \dots \wedge q_k^{(i)})$$

where

$$q_j^{(i)} = \begin{cases} q_j & \text{if } v_i(q_j) = T \\ \neg q_j & \text{otherwise} \end{cases} \quad \text{and} \quad \vec{q} \text{ is } q_1, \dots, q_k$$

Clearly from the definition of M and η , we have $\varphi(\vec{p}, \vec{q}) \models \eta(\vec{q})$. Let us now argue that $\eta(\vec{q}) \models \psi(\vec{q}, \vec{r})$. Consider an arbitrary valuation v such that $v[\psi] = F$. Since we have $\varphi \models \psi$, it must be that $v[\varphi] = F$. Consider any assignment v' such that $v \upharpoonright_{\vec{q}, \vec{r}} = v' \upharpoonright_{\vec{q}, \vec{r}}$. Since v and v' agree on the propositions appearing in ψ , we have $v'[\psi] = F$. Again, since ψ is a logical consequence of φ , it must be that $v'[\varphi] = F$. Thus, no matter how the assignment to propositions \vec{p} is changed from v , we will not be able to satisfy φ . This means that $v \upharpoonright_{\vec{q}} \notin M$. Thus, by our construction of η , $v[\eta] = F$. This establishes that $\eta \models \psi$. \square

Example 2.20 Let us look at a simple example that illustrates the construction of the interpolant in the proof of Theorem 2.19. Consider $\varphi = p \wedge (q_1 \vee q_2)$ and let $\psi = (q_1 \vee q_2 \vee r)$. It is easy to see that $\models \varphi \rightarrow \psi$. The set M constructed in the proof in this case would be $M = \{v_{TT}, v_{TF}, v_{FT}\}$, where v_{ij} is the function

$$v_{ij}(q_1) = i \quad \text{and} \quad v_{ij}(q_2) = j$$

Given M , the proof constructs the follow formula as interpolant.

$$\eta = (q_1 \wedge q_2) \vee (q_1 \wedge \neg q_2) \vee (\neg q_1 \wedge q_2).$$

2.3.2 Size of Interpolants

The interpolant η constructed in the proof of Theorem 2.19 is exponential in the number of common variables, and therefore, could be exponential in the size of the formulas φ and ψ . Can this be improved? Small interpolants can have a big impact in

the contexts where interpolants are used, like in formal verification. Or can we prove that, in the worst case, the interpolant needs to be exponential in the size of the input formulas? Unfortunately, like many questions in theoretical computer science, this remains open and unresolved — we cannot prove or disprove that the construction in the proof of Theorem 2.19 is the best. However, in this section, we will show that it is unlikely that we can construct polynomial sized interpolants for all formulas. Or more precisely, we will show that resolving whether there exist polynomial sized interpolants for all formulas, is closely related to other open questions in complexity theory.

In order to present these results on the size of interpolants, we need to introduce new circuit complexity classes. Circuit complexity for a problem is based on a *non-uniform* model of computation. The idea is the following. Imagine you have the ability to choose a different algorithm for each input length; how much resource would you need? The “programs” for each input length are circuits, and different aspects of these circuits correspond to different computational resources one may care about. Running time in this context, roughly corresponds to circuit size. This leads us to analogs of P , NP , $coNP$ in the context of non-uniform complexity, which are defined next.

Definition 2.21 (Circuits)

A *Boolean Circuit* C is a sequence of assignments A_1, A_2, \dots, A_n , where each A_i is one of the following forms.

$$\begin{aligned} P_i &= \mathsf{F} \\ P_i &= \mathsf{T} \\ P_i &= ? \\ P_i &= P_j \wedge P_k, \quad j, k < i \\ P_i &= P_j \vee P_k, \quad j, k < i \\ P_i &= \neg P_j, \quad j < i \end{aligned}$$

where each P_i is a variable that appears on the left-hand side in only A_i . The size of such a circuit, denoted $|C|$, is n .

The variable P_i is said to be an *input variable* if the line A_i corresponding to it is of the form $P_i = ?$. The input variables of C will be denoted by $I(C)$. Given an assignment $\nu : I(C) \rightarrow \{\mathsf{F}, \mathsf{T}\}$, the value of C under ν is the value assigned to the variable P_n in the last line. There is a natural order on the variables (based on which line they are assigned a value) which also imposes an order on the input variables. Thus, an assignment $I(C) \rightarrow \{\mathsf{F}, \mathsf{T}\}$ can be thought of as a string x , where $x[i]$ is the value assigned to the i th input variable. Under such an interpretation, the value of C under string x , will be denoted by $C(x)$.

Example 2.22 Circuits and formulas are two ways to represent Boolean functions. It is instructive to see how they differ by looking at an example.

Consider the boolean function $\text{parity}(x_1, x_2, \dots, x_n)$ which computes whether the number of propositions in $\{x_1, x_2, \dots, x_n\}$ set to T is odd or even. For example, we could write down the formula for some simple cases.

$$\begin{aligned}\text{parity}(x_1, x_2) &= (x_1 \vee x_2) \wedge \neg(x_1 \wedge x_2) \\ \text{parity}(x_1, x_2, x_3) &= (((x_1 \vee x_2) \wedge \neg(x_1 \wedge x_2)) \vee x_3) \\ &\quad \wedge \neg(((x_1 \vee x_2) \wedge \neg(x_1 \wedge x_2)) \wedge x_3)\end{aligned}$$

More generally, we can inductively write down the formula for $\text{parity}(x_1, x_2, \dots, x_n)$ by observing that the parity of x_1, x_2, \dots, x_n is odd iff either (a) the parity of x_1, \dots, x_{n-1} is even and x_n is T , or (b) parity of x_1, \dots, x_{n-1} is odd and x_n is F . This results in the following definition.

$$\begin{aligned}\text{parity}(x_1, x_2, \dots, x_{n-1}, x_n) &= (\text{parity}(x_1, x_2, \dots, x_{n-1}) \vee x_n) \\ &\quad \wedge \neg(\text{parity}(x_1, x_2, \dots, x_{n-1}) \wedge x_n)\end{aligned}$$

Observe that the formula we wrote down for $\text{parity}(x_1, x_2, x_3)$ is based on exactly this definition. Observe that the size of the formula $\text{parity}(x_1, \dots, x_n)$ is double the size of $\text{parity}(x_1, \dots, x_{n-1})$. This means that the size of parity for n arguments is $O(2^n)$.

A circuit for the same formula does not grow as rapidly. This is because it can “reuse” previous computed answers. Let us look at the circuit corresponding to the formula $\text{parity}(x_1, x_2, x_3)$ we wrote above.

$$\begin{aligned}x_1 &= ? \\ x_2 &= ? \\ x_3 &= ? \\ P_4 &= x_1 \vee x_2 \\ P_5 &= x_1 \wedge x_2 \\ P_6 &= \neg P_5 \\ P_7 &= P_4 \wedge P_6 \\ P_8 &= P_7 \vee x_3 \\ P_9 &= P_7 \wedge x_3 \\ P_{10} &= \neg P_9 \\ P_{11} &= P_8 \wedge P_{10}\end{aligned}$$

Notice that variable P_7 stores $\text{parity}(x_1, x_2)$ and it is reused without paying an extra cost. More generally, if C_{n-1} is the circuit computing $\text{parity}(x_1, \dots, x_{n-1})$ with last line $P_{k_{n-1}}$, the circuit C_n computing $\text{parity}(x_1, \dots, x_n)$ is given by

$$\begin{aligned}C_{n-1} \\ x_n &= ? \\ Q_1 &= P_{k_{n-1}} \vee x_n \\ Q_2 &= P_{k_{n-1}} \wedge x_n \\ Q_3 &= \neg Q_2 \\ P_{k_n} &= Q_1 \wedge Q_3\end{aligned}$$

Now if $|C_n| = |C_{n-1}| + 5$. Therefore, in general, we get $|C_n| = O(n)$. Thus the circuit for the same syntactic formula can be exponentially smaller.

Definition 2.23 A language $A \subseteq \{0, 1\}^*$ is said to be in $\mathsf{P/poly}$ iff there are constants $c, k \in \mathbb{N}$, and a (infinite) family of circuits $\{C_i\}_{i \in \mathbb{N}}$ such that (a) for every n , $|C_n| \leq cn^k$, and (b) for every x , $x \in A$ iff $C_{|x|}(x) = \top$.

A language $A \subseteq \{0, 1\}^*$ is said to be in $\mathsf{NP/poly} (\mathsf{coNP/poly})$ iff there are constants $c, k \in \mathbb{N}$, and a (infinite) family of circuits $\{C_i\}_{i \in \mathbb{N}}$ such that (a) for every n , $|C_n| \leq cn^k$, and (b) for every x , $x \in A$ iff for *some* (*all*) r , $C_{|x|}(x, r) = \top$.

One can think of $\mathsf{P/poly}$ (or $\mathsf{NP/poly}$ or $\mathsf{coNP/poly}$) as the collection of problems that be solved in polynomial time (or nondeterministic polynomial time or co-nondeterministic polynomial time) *given* a polynomially long *advice string*, namely, the description of the appropriate circuit. Just like in the case of (uniform/regular) complexity classes where it is open whether $\mathsf{P} \stackrel{?}{=} \mathsf{NP} \cap \mathsf{coNP}$, the same question is open in the non-uniform context as well. That is, a long standing open problem is whether $\mathsf{P/poly} \stackrel{?}{=} \mathsf{NP/poly} \cap \mathsf{coNP/poly}$.

Mundici's theorem says that establishing the existence of interpolants with polynomial circuit representation for all pairs of formulas, is equivalent to resolving the $\mathsf{P/poly}$ versus $\mathsf{NP/poly} \cap \mathsf{coNP/poly}$ question. Thus, it is likely to be difficult to establish.

Theorem 2.24 (Mundici)

If for any φ and ψ such that $\varphi \models \psi$ there is an interpolant whose circuit size is polynomial in φ and ψ then

$$\mathsf{P/poly} = \mathsf{NP/poly} \cap \mathsf{coNP/poly}$$

Proof Consider a problem $L \in \mathsf{NP/poly} \cap \mathsf{coNP/poly}$. Thus, there are families of circuits $\{A_n\}_{n \in \mathbb{N}}$ and $\{B_n\}_{n \in \mathbb{N}}$ such that for each n , A_n and B_n have size bounded by a polynomial function of n and for any binary string w , $w \in L$ iff $\exists p. A_{|w|}(p, w) = \top$ iff $\forall r. B_{|w|}(w, r) = \top$. These observations are just a consequence of $L \in \mathsf{NP/poly}$ and $L \in \mathsf{coNP/poly}$.

Based on the previous paragraph, we have, for any n , if for some assignment of values to \vec{q} , if $\exists \vec{p}. A_n(\vec{p}, \vec{q})$ holds then it also the case that $\forall \vec{r}. B_n(\vec{q}, \vec{r})$ holds. Therefore, we have $A_n(\vec{p}, \vec{q}) \models B_n(\vec{q}, \vec{r})$. By our assumption on interpolants, we have an interpolant (as a circuit) $C_n(\vec{q})$ such that $|C_n|$ is bounded by a polynomial in $|A_n|$ and $|B_n|$. Since $|A_n|$ and $|B_n|$ are bounded by polynomials in n , we have $|C_n|$ is bounded by a polynomial in n . Further, since C_n is an interpolant, we have if $\exists \vec{p}. A_n(\vec{p}, \vec{q})$ holds then $C_n(\vec{q})$ holds, and if $C_n(\vec{q})$ holds then $\forall \vec{r}. B_n(\vec{q}, \vec{r})$ holds. Thus, $\{C_n\}$ is a family of polynomially sized circuits deciding L , and thereby demonstrating that $L \in \mathsf{P/poly}$. \square

2.3.3 Interpolants from Refutations

Constructing small interpolants for all formulas is likely to be difficult. However, it turns out that, if a formula $A(\vec{p}, \vec{q}) \rightarrow B(\vec{q}, \vec{r})$ has a short proof in some

proof systems, then the proof can be used to construct an interpolant, whose size is propositional to the original proof. One proof system that admits such a result is resolution.

Theorem 2.25 *Let the collection of clauses $\Gamma = \{A_i(\vec{p}, \vec{q})\}_{i=1}^k \cup \{B_j(\vec{q}, \vec{r})\}_{j=1}^\ell$ have a resolution refutation of length n . Then there is a circuit $C(\vec{q})$ such that*

$$\bigwedge_i A_i(\vec{p}, \vec{q}) \models C(\vec{q}) \text{ and } C(\vec{q}) \wedge \bigwedge_j B_j(\vec{q}, \vec{r}) \text{ is unsatisfiable.}$$

Further $|C|$ is $O(n)$.

Proof Since $\Gamma = \{A_i(\vec{p}, \vec{q})\}_i \cup \{B_j(\vec{q}, \vec{r})\}_j$ is unsatisfiable, every truth valuation v fails to satisfy at least one of the $A_i(\vec{p}, \vec{q})$ or $B_j(\vec{q}, \vec{r})$. We can also restate the properties of an interpolant C as

$$\begin{aligned} \models \neg C(\vec{q}) &\rightarrow \neg \bigwedge_i A_i(\vec{p}, \vec{q}) \text{ and} \\ \models C(\vec{q}) &\rightarrow \neg \bigwedge_j B_j(\vec{q}, \vec{r}) \end{aligned}$$

Thus, $C(\vec{q})$ can be thought of as a way of labelling truth valuations: those labelled F will not satisfy some clause $A_i(\vec{p}, \vec{q})$ and those labelled T will not satisfy some $B_j(\vec{q}, \vec{r})$.

The structure of the proof will be as follows. Let $\psi_1, \psi_2, \dots, \psi_n$ be a resolution refutation of Γ . Let the set of common variables $\vec{q} = \{q_1, \dots, q_m\}$. Our circuit for the interpolant will be a sequence of the form $q_1 = ?, q_2 = ?, \dots, q_m = ?, P_1 = E_1, P_2 = E_2, \dots, P_n = E_n$ — the first m lines asserting that q_i 's are variables, and then having one line $P_i = E_i$ corresponding to each line ψ_i of our refutation. It will be convenient to consider expressions E_i on the right hand side that have more than one logical connective. We will find it convenient to talk about the “circuit corresponding to a line ψ_i in the refutation”. What we mean by this is look at the sequence of assignments upto (and including) line $P_i = E_i$; we will denote this as circuit C_i . The value of C_i (with respect to a truth assignment) will simply be the value variable P_i gets when we compute this circuit.

As mentioned above, we will construct the circuit line by line, corresponding to the refutation. With each line ψ_i in the refutation, we can associate a set of truth assignments, namely those that *do not* satisfy ψ_i , i.e., $M_i = \{v \mid v[\psi_i] = \mathsf{F}\}$. The invariant we will maintain as we build the circuit line by line, is that C_i “correctly labels” assignments belonging to M_i . That is, for $v \in M_i$, if $C_i(v) = \mathsf{F}$ then v does not satisfy some clause $A_i(\vec{p}, \vec{q})$ and if $C_i(v) = \mathsf{T}$ then v does not satisfy some clause $B_j(\vec{q}, \vec{r})$. Notice that the last line $\psi_n = \{\}$, and so M_n is the set of all truth assignments. Thus, if the invariant is maintained, C_n will indeed be an interpolant because it will “correctly” label all assignments.

Let us now describe how we construct the circuit. For each line ψ_e , we will add a line $P_e = E_e$. What E_e is will depend on the justification for the line ψ_e in the refutation. Let us begin with the cases when ψ_e is a clause belonging to Γ . There are two cases to consider.

- Suppose $\psi_e \in \{A_i(\vec{p}, \vec{q})\}_{i=1}^k$. Observe that any $v \in M_e$ does not satisfy a clause in $\{A_i(\vec{p}, \vec{q})\}_{i=1}^k$ and could be safely labeled F . Thus, the line corresponding to ψ_e will be $P_e = F$.
- If $\psi_e \in \{B_j(\vec{q}, \vec{r})\}_{j=1}^\ell$, then each $v \in M_e$ could be labeled T because it does not satisfy $\psi_e \in \{B_j(\vec{q}, \vec{r})\}_{j=1}^\ell$. Thus, we have $P_e = T$.

The next cases to consider are when ψ_e is a resolvent of two clauses. So let $\psi_e = \rho_1 \cup \rho_2$ and let it be the resolvent of clauses ψ_a and ψ_b in the refutation. We need to consider different cases based on the proposition with respect to which ψ_e is a resolvent. Let us begin by considering the case when the proposition being resolved is one of the common variables. Without loss of generality, let us take $\psi_a = \rho_1 \cup \{q\}$ and $\psi_b = \rho_2 \cup \{\neg q\}$. Consider an arbitrary assignment $v \in M_e$, i.e., $v[\psi_e] = F$. From the soundness of the resolution proof rule, we know that either $v \in M_a$ or $v \in M_b$. If $v(q) = F$ then $v \in M_a$; it may also, in addition, be the case that $v \in M_b$, but that is unimportant. We could label such assignments in the same manner as the labeling corresponding to line ψ_a , i.e., as per the value of variable P_a . Similarly, if $v(q) = T$ then $v \in M_b$ and so it can be labeled in the same manner as P_b . This gives us that the line corresponding to ψ_e in this case should be $P_e = (\neg q \wedge P_a) \vee (q \wedge P_b)$.

Let us now consider the case when ψ_e is a resolvent of ψ_a and ψ_b , but the resolution step is taken with respect to a proposition (say s) that is either in \vec{p} or in \vec{r} . Once again any $v \in M_e$ must also belong to $M_a \cup M_b$, but now since s is not in \vec{q} it cannot be explicitly mentioned in the line $P_e = E_e$, as we did in the previous case. Again, without loss of generality, let us assume that $\psi_e = \rho_1 \cup \rho_2$, $\psi_a = \rho_1 \cup \{s\}$, $\psi_b = \rho_2 \cup \{\neg s\}$. Let v' be the assignment that is identical to v , except that it flips the assignment to s . Without loss of generality, we can assume that $v(s) = F$ and $v'(s) = T$, and for all other propositions t , $v(t) = v'(t)$. Observe that $v \in M_a$ and $v' \in M_b$. Further, $C_a(v) = C_a(v')$ and $C_b(v) = C_b(v')$; this is because C_a and C_b do not mention s . Let us consider two cases based on whether $s \in \vec{p}$ or $s \in \vec{r}$.

- Suppose $s \in \vec{p}$. Observe that in this case, for any j , $v[B_j] = v'[B_j]$, because s does not appear in B_j . Now, since C_a and C_b satisfy our invariant, we have, if $C_a(v) = T (= C_a(v'))$ then for some j , $v[B_j] = F (= v'[B_j])$. Similarly, if $C_b(v') = T (= C_b(v))$ then for some j , $v'[B_j] = F (= v[B_j])$. Thus, if either C_a or C_b label v, v' by T , then C_e must do the same. Otherwise, both C_a and C_b label v and v' as F , and this common label is correct as per our invariant. Therefore, we have $P_e = P_a \vee P_b$ in this case.
- Now let us consider $s \in \vec{q}$. In this case, for any i , $v[A_i] = v'[A_i]$. Using an argument similar to the previous case, we can argue that if either C_a or C_b label v, v' by F then C_e must do the same. Otherwise, C_a and C_b agree on the label, and that is indeed the correct label. Therefore, dually, in this case we have, $P_e = P_a \wedge P_b$.

The proof that the invariant is maintained follows inductively, from the arguments we have made for each case. Thus, C_n is indeed the interpolant. It is worth noting that in C_n we have a sequence of initial assignments $q_1 = ?, q_2 = ?, \dots, q_m = ?$ that assert that \vec{q} are input variables. This seems to suggest that the size of C_n also depends on

$|\vec{q}|$ and hence on Γ . However, instead of asserting that all variables in \vec{q} are input variables, we could assert only those variables in \vec{q} that appear in the refutation. Thus, $|C_n|$ is indeed linear in the size of the refutation. \square

Example 2.26 Let us look at an example to see how an interpolant can be constructed from a refutation. Consider the set of clauses

$$\Gamma = \overbrace{\{\{p, q\}, \{\neg p, r\}\}}^A, \overbrace{\{\{\neg q, r\}, \{\neg r\}\}}^B.$$

Here p is a proposition that only appears in the A -clauses, and q and r are propositions that appear in both A and B clauses. Γ is unsatisfiable, and we are looking for an interpolant that only mentions the common variables q, r . Below we have the refutation (on the left) alongside the circuit for the interpolant (on the right) as per the proof of Theorem 2.25.

$$\begin{array}{ll} q = ? & \\ r = ? & \\ \{\neg p, r\} & P_1 = 0 \\ \{\neg r\} & P_2 = 1 \\ \{\neg p\} & P_3 = (\neg r \wedge P_1) \vee (r \wedge P_2) \\ \{\neg q, r\} & P_4 = 1 \\ \{\neg q\} & P_5 = (\neg r \wedge P_4) \vee (r \wedge P_2) \\ \{p, q\} & P_6 = 0 \\ \{q\} & P_7 = P_3 \vee P_6 \\ \{\} & P_8 = (\neg q \wedge P_7) \vee (q \wedge P_5) \end{array}$$

Observations like Theorem 2.25 have also been established for other proof systems of propositional logic. That is, in these proof systems, a short proof for a fact can be converted into a construction of a small interpolant. Theorem 2.25 can be strengthened for certain special sets of clauses — one can show that in certain special cases, not only is the interpolant small, but it is also *monotonic*.

Definition 2.27 (Monotone Circuits)

A *monotone circuit* C is one where there are no assignments of the form $P_i = \neg P_j$.

A monotone circuit enjoys the following monotonic property. Let us say $v_1 \leq v_2$ if for all proposition p , if $v_1(p) = T$ then $v_2(p) = T$, i.e., v_2 sets at least as many propositions to T as v . The value of a monotone circuit with respect to this ordering on assignments is monotonic. In other words, if C is monotone, then for any v_1, v_2 such that $v_1 \leq v_2$, we have $C(v_1) = T$ implies $C(v_2) = T$.

Theorem 2.28 *Let the collection of clauses $\Gamma = \{A_i(\vec{p}, \vec{q})\}_{i=1}^k \cup \{B_j(\vec{q}, \vec{r})\}_{j=1}^\ell$ have a resolution refutation of length n . Further assume that either \vec{q} occur only positively in A_i s or \vec{q} occur only negatively in B_j s. Then there is a monotone circuit $C(\vec{q})$ such that*

$$\bigwedge_i A_i(\vec{p}, \vec{q}) \models C(\vec{q}) \text{ and } C(\vec{q}) \wedge \bigwedge_j B_j(\vec{q}, \vec{r}) \text{ is unsatisfiable.}$$

In addition, $|C|$ is $O(n)$.

Proof The construction of the interpolant is identical to that in the proof of Theorem 2.25. All cases in that proof go through except for the case when the line ψ_e is a resolvent with respect to proposition $q \in \vec{q}$, i.e., the common variables. In this case, in the proof of Theorem 2.25, the circuit was $P_e = (\neg q \wedge P_a) \vee (q \wedge P_b)$, where ψ_e is the resolvent of lines $\psi_a = \rho_1 \cup \{p\}$ and $\psi_b = \rho_2 \cup \{\neg q\}$. This is not monotonic because if the use of $\neg q$. To prove our result, we need to change the circuit in this case. We will change it by forcing it to be monotone in the most naïve way — we will remove the offending $\neg q$ and write $P_e = P_a \vee (q \wedge P_b)$.

The resulting construction is correct, but the inductive argument using the invariant from Theorem 2.25 does not go through! Let us see what the problem is. Recall, that for any line ψ_e , we defined the set $M_e = \{v \mid v[\psi_e] = F\}$. The invariant we proved in Theorem 2.25 was for any valuation $v \in M_e$, we have

- If $v[C_e] = F$ then $v[A_i] = F$ for some i , and
- If $v[C_e] = T$ then $v[B_j] = F$ for some j .

Let us try to prove this invariant by induction as in the proof of Theorem 2.25. Consider the case that we just changed, i.e., of a resolvent with respect to a common variable. So $\psi_e = \rho_1 \cup \rho_2$ is the resolvent of lines $\psi_a = \rho_1 \cup \{q\}$ and $\psi_b = \rho_2 \cup \{\neg q\}$. And we have, $P_e = P_a \vee (q \wedge P_b)$. Consider a valuation $v \in M_e$. The problem with carrying out this inductive proof occurs when $v(q) = T$; the other case of $v(q) = F$ goes through rather simply. Then $v \in M_b$. Now if $C_a(v) = F$ or $C_b(v) = T$ then $C_e(v) = C_b(v)$ and correctness follows from the inductive assumptions on C_b . The problem occurs when $C_a(v) = T$ and $C_b(v) = F$.

The way to fix the problem is to prove a *stronger* invariant. In our old invariant, we proved that our circuit for line e was correct on the truth assignments in $M_e = \{v \mid v[\psi_e] = F\}$. Our strengthening will show that the circuit for line e is correct on a larger set of truth assignments. Depending on whether we consider the case when \vec{q} appears only positively in $\{A_i(\vec{p}, \vec{q})\}_i$ or the case when \vec{q} occurs only negatively in $\{B_j(\vec{q}, \vec{r})\}_j$, the invariant (and the proof) is slightly different. We will present the proof only for the case when \vec{q} occurs negatively in $\{B_j(\vec{q}, \vec{r})\}_j$. We will state the modified invariant for the other case, but leave the details to be filled out by the reader.

To describe the invariant in the case when \vec{q} occurs negatively in $\{B_j(\vec{q}, \vec{r})\}_j$, we need to introduce some notation. For a clause ψ , $\psi \upharpoonright_{\vec{p}, \vec{q}}$ be the clause obtained by removing all literals involving propositions in \vec{r} . On the other hand, $\psi \upharpoonright_{\neg \vec{q}, \vec{r}}$ is the clause obtained from ψ by removing all literals of proposition \vec{p} as well as all positive literals of \vec{q} . Our stronger invariant will be for every valuation v

- if $v[\psi \upharpoonright_{\vec{p}, \vec{q}}] = F$ and $C_e(v) = F$ then $v[A_i] = F$ for some i , and
- if $v[\psi \upharpoonright_{\neg \vec{q}, \vec{r}}] = F$ and $C_e(v) = T$ then $v[B_j] = F$ for some j .

Notice that since $\{\} \upharpoonright_{\vec{p}, \vec{q}} = \{\} \upharpoonright_{\neg \vec{q}, \vec{r}} = \{\}$, proving this new invariant guarantees that C_n (where n is length of the resolution refutation) is an interpolant.

We now argue that the stronger invariant holds for the new construction. We consider each case in order.

$\psi_e \in \{A_i\}_i$: In this case, we have $\psi_e \upharpoonright_{\vec{p}, \vec{q}} = \psi_e$ and $P_e = \mathsf{F}$. The invariant, therefore, holds.

$\psi_e \in \{B_j\}_j$: Again we have $\psi_e \upharpoonright_{\vec{q}, \vec{r}} = \psi_e$. Since $P_e = \mathsf{T}$, the invariant holds.

Resolvent w.r.t \vec{q} : Let $\psi_e = \rho_1 \cup \rho_2$ be the resolvent of lines $\psi_a = \rho_1 \cup \{q\}$ and $\psi_b = \rho_2 \cup \{\neg q\}$. Recall we have $P_e = P_a \vee (q \wedge P_b)$.

1. Consider v such that $v[\psi_e \upharpoonright_{\vec{p}, \vec{q}}] = \mathsf{F}$ and $C_e(v) = \mathsf{F}$. In this case, if $v(q) = \mathsf{T}$ then $v[\psi_b \upharpoonright_{\vec{p}, \vec{q}}] = v[\psi_b] = \mathsf{F}$. Also, since $C_e(v) = \mathsf{F}$, it must be that $C_b(v) = \mathsf{F}$, and so correctness follows by induction. On the other hand, if $v(q) = \mathsf{F}$ then $v[\psi_a \upharpoonright_{\vec{p}, \vec{q}}] = v[\psi_a] = \mathsf{F}$. Also, $C_a(v) = \mathsf{F}$ and correctness follows by induction.
2. Consider v such that $v[\psi_e \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$ and $C_e(v) = \mathsf{T}$. If $C_a(v) = \mathsf{T}$ then since $v[\psi_a \upharpoonright_{\vec{q}, \vec{r}}] = v[\rho_1 \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$, the invariant follows by induction. Notice, how the stronger invariant helped the proof go through in this case which was problematic before. On the other hand, if $v[q \wedge C_b] = \mathsf{T}$ then $v(q) = \mathsf{T}$. So $v[\psi_b \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$ and then invariant holds by induction.

Resolvent w.r.t \vec{p} : Let $\psi_e = \rho_1 \cup \rho_2$ be the resolvent of lines $\psi_a = \rho_1 \cup \{p\}$ and $\psi_b = \rho_2 \cup \{\neg p\}$. Recall $P_e = P_a \vee P_b$.

1. Suppose $C_e(v) = \mathsf{F}$ and $v[\psi_e \upharpoonright_{\vec{p}, \vec{q}}] = \mathsf{F}$. Then we know $C_a(v) = C_b(v) = \mathsf{F}$. Further either $v[\psi_a] = \mathsf{F}$ or $v[\psi_b] = \mathsf{F}$. Thus, correctness of construction by induction.
2. Suppose $v[\psi_e \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$ and $C_e(v) = \mathsf{T}$. Now $\psi_e \upharpoonright_{\vec{q}, \vec{r}} = \psi_a \upharpoonright_{\vec{q}, \vec{r}} \cup \psi_b \upharpoonright_{\vec{q}, \vec{r}}$, and so $v[\psi_a \upharpoonright_{\vec{q}, \vec{r}}] = v[\psi_b \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$. Further since $C_e(v) = \mathsf{T}$, either $C_a(v) = \mathsf{T}$ or $C_b(v) = \mathsf{T}$. So correctness follows by induction.

Resolvent w.r.t \vec{r} : Proof similar to previous case.

The proof of correctness when \vec{q} appears positively in $\{A_i(\vec{p}, \vec{q})\}_i$ is similar, though the invariant is slightly different. For a clause ψ , take $\psi \upharpoonright_{\vec{p}, \vec{q}}$ to be the clause obtained by removing literals of \vec{r} and negative literals of \vec{q} . In addition, $\psi \upharpoonright_{\vec{q}, \vec{r}}$ is the clause obtained by removing literals of \vec{p} . The invariant we will prove about the construction is, for every v ,

- if $v[\psi \upharpoonright_{\vec{p}, \vec{q}}] = \mathsf{F}$ and $C_e(v) = \mathsf{F}$ then $v[A_i] = \mathsf{F}$ for some i , and
- if $v[\psi \upharpoonright_{\vec{q}, \vec{r}}] = \mathsf{F}$ and $C_e(v) = \mathsf{T}$ then $v[B_j] = \mathsf{F}$ for some j .

The proof is similar and skipped. \square

Example 2.29 Let us consider the set of clauses Γ from Example 2.26.

$$\Gamma = \overbrace{\{\{p, q\}, \{\neg p, r\}\}}^A, \overbrace{\{\neg q, r\}, \{\neg r\}\}}^B.$$

Notice that the common propositions, q and r , appear only positively in A . The refutation alongside the interpolant construction is as follows.

$$\begin{array}{ll}
 q = ? & \\
 r = ? & \\
 \{\neg p, r\} & P_1 = 0 \\
 \{\neg r\} & P_2 = 1 \\
 \{\neg p\} & P_3 = P_1 \vee (r \wedge P_2) \\
 \{\neg q, r\} & P_4 = 1 \\
 \{\neg q\} & P_5 = P_4 \vee (r \wedge P_2) \\
 \{p, q\} & P_6 = 0 \\
 \{q\} & P_7 = P_3 \vee P_6 \\
 \{\} & P_8 = P_7 \vee (q \wedge P_5)
 \end{array}$$

2.3.4 Lower bounds on Resolution Refutations

The results in Sect. 2.2 connecting resolution refutation lengths and size of interpolants, allows one to prove lower bounds on the length of resolution refutations for formulas. In particular we can show that there are sets of clauses Γ for which the shortest resolution refutations are exponential in the size of Γ . Thus, not every unsatisfiable formula has a short proof in resolution. The specific example we consider relates to cliques in graphs and their coloring. Let us recall these classical problems.

Recall that in Proposition 1.35, we showed that determining if a graph is k -colorable can be reduced to checking the satisfiability of a set of formulas. More specifically, let us fix the graph $G = (V, E)$ to have n vertices. Any such graph can be represented by an assignment to propositions $\{q_{uv} \mid u, v \in \{1, 2, \dots, n\}\}$, with the interpretation that $(u, v) \in E$ iff q_{uv} is set to \top . Using r_{ui} to denote “vertex u has color i ”, there is a set of clauses $\text{color}_{n,k}(\vec{q}, \vec{r})$ such that $\vec{v} \models \text{color}_{n,k}(\vec{q}, \vec{r})$ iff the graph (over n vertices) represented by $\vec{v} \upharpoonright_{\vec{q}}$ has a k -coloring given by $\vec{v} \upharpoonright_{\vec{r}}$. The set $\text{color}_{n,k}$ is almost identical to the construction given in the proof of Proposition 1.35. The only difference is that we will use the clause $\neg q_{uv} \vee \neg r_{ui} \vee \neg r_{vi}$, for every pair of vertices $u, v \in \{1, 2, \dots, n\}$ and color $i \in \{1, 2, \dots, k\}$, instead of $\neg r_{ui} \vee \neg r_{vi}$ for every edge (u, v) as given in Proposition 1.35. Note that $\text{color}_{n,k}$ has $O(n^2k + nk^2)$ clauses.

Whether a graph is k -colorable is related to the presence of graph structures called *cliques*.

Definition 2.30 A k -clique in a graph $G = (V, E)$ is a subset $U \subseteq V$ such that $|U| = k$ and for every $u, v \in U$, with $u \neq v$, we have $(u, v) \in E$.

Like graph coloring, the problem of determining if a graph has a k -clique can be reduced to satisfiability.

Proposition 2.31 For any n, k , there is a set of $O(n^2 k^2)$ clauses $\text{clique}_{n,k}(\vec{p}, \vec{q})$ such that $\mathbf{v} \models \text{clique}_{n,k}(\vec{p}, \vec{q})$ iff the graph represented by $\mathbf{v} \upharpoonright_{\vec{q}}$ has a k -clique given by $\mathbf{v} \upharpoonright_{\vec{p}}$

Proof The proof of this observation is similar to that of Proposition 1.35. We will introduce propositions that encode the k -clique, and clauses will specify constraints that characterize properties of a k -clique. Let proposition p_{iu} , for $i \in \{1, \dots, k\}$ and $u \in \{1, \dots, n\}$, denote that “the i th vertex in clique is u ”. Then $\text{clique}_{n,k}(\vec{p}, \vec{q})$ is the following set of clauses.

- For each $1 \leq i \leq k$, the clause $p_{i1} \vee p_{i2} \vee \dots \vee p_{in}$. These clauses capture the constraint that the i th vertex of the clique must be among $\{1, \dots, n\}$.
- For each $1 \leq i \leq k$, and $1 \leq u, v \leq n$ such that $u \neq v$, the clause $\neg p_{iu} \vee \neg p_{iv}$. Intuitively, this says that the i th vertex of the clique can be at most one vertex.
- For each $1 \leq i, j \leq k$ with $i \neq j$, and $1 \leq u \leq n$, the clause $\neg p_{iu} \vee \neg p_{ju}$. These clauses say that the i th and j th vertex of the clique cannot be the same vertex u .
- For each $1 \leq i, j \leq k$ and $1 \leq u, v \leq n$ with $u \neq v$, we have the clause $\neg p_{iu} \vee \neg p_{ju} \vee q_{uv}$. These clauses together say that if u, v are vertices in the clique then they have an edge between them.

The proof that these clauses satisfy the proposition is left to the reader. \square

Observe that if a graph G has a k -clique then it cannot be colored using $k - 1$ colors. This is because each of the vertices in the clique must get different colors. Thus a graph with a k -clique needs at least k colors. This leads us to the following observation.

Proposition 2.32 For any n, k , $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ is unsatisfiable.

Proof A satisfying assignment for $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ is a graph encoded by \vec{q} that has k -clique (identified by \vec{p}) and can be colored using $k - 1$ colors (with the coloring encoded by \vec{r}). This is clearly impossible. \square

Since $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ is unsatisfiable, it must have a resolution refutation. How long is its refutation? Our goal will be to prove that this is exponential in n . Since the size of $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ itself is polynomial in n , this would be an example that has a “long” proof. In order to establish this result, we present a celebrated result in circuit complexity whose proof can be found in textbooks like [?].

Theorem 2.33 (Razborov, Alon-Bopanna)

Any monotone circuit that evaluates to T on n -vertex graphs containing a k -clique and evaluates to F on n -vertex graphs that are $k - 1$ colorable must have size at least $n^{\Omega(\sqrt{k})}$, when $k \leq n^{\frac{1}{4}}$.

Theorem 2.33 combined with Theorem 2.28 gives us the desired lower bound on proof lengths.

Theorem 2.34 Any resolution refutation of $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ must have length at least $n^{\Omega(\sqrt{k})}$, when $k \leq n^{\frac{1}{4}}$.

Proof Let there be a resolution refutation of length ℓ . Observe that \vec{q} appears only positively in $\text{clique}_{n,k}(\vec{p}, \vec{q})$ and only negatively in $\text{color}_{n,k}(\vec{q}, \vec{r})$. Thus, $\text{clique}_{n,k}(\vec{p}, \vec{q}) \cup \text{color}_{n,k-1}(\vec{q}, \vec{r})$ satisfies the conditions of Theorem 2.28, and so there is a monotone interpolant of size $O(\ell)$. The interpolant is a monotone circuit satisfying conditions of Theorem 2.33. Thus, ℓ must be at least $n^{\Omega(\sqrt{k})}$. \square

Theorem 2.34 establishes that resolution proofs can be long as an exponential function of the size of the input clauses. Historically, the above theorem was not the first example of a proof that some formulas have long resolution proofs. The first such result was established by Haken. He showed that the *pigeon hole principle* has long resolution proofs. Recall that the pigeon hole principle says that if there are n holes and $n + 1$ pigeons then some hole may contain more than one pigeon. Let pigeonhole_n be the propositional logic formula that says that every pigeon goes to a hole and no hole contains more than one pigeon. Then pigeonhole_n is unsatisfiable, and Haken showed that any resolution refutation of this fact must be exponential in n . The broad principle of using interpolation and lower bounds from monotonic circuit complexity have been used to establish that other proof systems can also have long proofs.

Understanding how long resolution proofs can be, and when they can be long, helps our theoretical understanding of the limits of SAT solvers — what examples they may work well and when they can take long. But a probably more important reason for studying proof lengths in some proof system is because of its connections to some fundamental questions in complexity theory. Essentially, the goal in proof complexity is to understand if there is a proof system for propositional logic that has the property that all facts have short proofs. Investigating whether this is true or not has important implications in complexity theory. Let us see why.

Definition 2.35 A proof system Π is *super* if every tautology φ has a proof in Π such that length of the proof is bounded by a polynomial function of $|\varphi|$.

Now Theorem 2.34 says that resolution is not a super proof system. But are there other proof systems that are super? This is intimately tied to another open question in complexity theory.

Theorem 2.36 (Cook-Reckhow)

Propositional logic has a super proof systems if and only if $\text{NP} = \text{coNP}$.

Proof There are two directions to this proof. Assume that there is a super proof system. Then the problem to determine if a given formula is valid, is in NP — the NP algorithm simply guesses the proof and checks that it is a proof in our super proof system. Since the problem of checking validity is coNP -complete, it follows that $\text{coNP} = \text{NP}$; the NP algorithm for an arbitrary problem $A \in \text{coNP}$ is simple to reduce it to validity checking and then use the NP algorithm based on the super proof system.

On the other hand, suppose $\text{NP} = \text{coNP}$ then there is nondeterministic Turing machine N , running in polynomial time, that checks if a given propositional logic formula is valid. Proofs for a formula φ in our new “super” proof system will simply

be the nondeterministic choices that cause N to accept φ ; notice, that these proofs will be polynomially long because N only has computations that are polynomially long. \square

In the light of Theorem 2.36, to resolve the NP versus coNP question, we need to prove that there are no super proof systems for proposition logic. Cook proposed an approach to tackling this problem. Consider concrete natural proof systems for propositional logic, one by one, and show that they are not super. Then use the intuition developed in this process to generalize and prove the absence of any super proof system. Resolution was the first proof system shown to be not super. Since then other proof systems have also been proved to be not super. However, there are still natural proof systems for which we have not been able to prove exponential lower bounds for proof lengths. One such proof system is the Frege proof system we introduced. It is still open whether there are tautologies for which proofs in the Frege proof system will be exponentially long.

Chapter 3

First Order Logic

Syntax, Semantics, and Overview

First order logic is a formal language to describe and reason about *predicates*. Modern efforts to study this logic grew out of a desire to study the foundations of mathematics in number theory and set theory. It has a careful treatment of functions, variables, and quantification. First order logic deals with predicates as opposed to propositions — declarative statements that are either true or false — which is the subject of study in propositional logic. A predicate is a proposition that depends on the value of some variables. For example truth of the statement “ x is prime”, depends on the value x takes (and of course also on the meaning of “is prime”). If $x = 2$ the statement “ x is prime” would be true and if $x = 4$ it would be false. Predicates may depend on more than one variable. For example, the truth of $P(x, y) \triangleq x + y = 0$ depends on the values of both x and y .

One way to convert a predicate into a proposition by substituting values for the predicate variables. For example, for the predicate P defined in the previous paragraph, $P(2, -2)$ denotes the proposition “ $2 + (-2) = 0$ ”. Another way to obtain propositions in predicate logic is by using *quantifiers*, which allows one to express statements like the predicate holds for all values of the variable, or the predicate holds for some values of the variable. In this chapter, we introduce the syntax and semantics of first order logic, and some of the questions we will explore in this book.

3.1 Syntax

First order logic formulas are defined over a *vocabulary* or *signature* that identifies *non-logical symbols*, namely, the predicates, constants, and functions that can be used in the formulas.

Definition 3.1 A *vocabulary* or *signature* is $\tau = \{C, \mathcal{F}, \mathcal{R}\}$, where

- $C = \{c_1, c_2, \dots\}$ is a set of *constant* symbols,
- $\mathcal{F} = \{\mathcal{F}^k\}_k$ is a collection of sets with $\mathcal{F}^k = \{f_1^k, f_2^k, \dots\}$ being the set of k -ary function symbols, and

- $\mathcal{R} = \{\mathcal{R}^k\}_k$ is a collection of sets with $\mathcal{R}^k = \{R_1^k, R_2^k, \dots\}$ being the set of k -ary relation symbols.

Note that any of the above sets of constants, k -ary function symbols or k -ary relation symbols can be empty, finite, or infinite. A signature is *purely relational* or simply *relational* if there are no constants or functions in the signature, i.e., $C = \mathcal{F} = \emptyset$. A signature is *finite* if the total number of symbols in the signature is finite.

We will typically consider signatures that are finite. When the arity of a function or relation symbol is clear from the context, we will drop the superscript.

Formulas in first-order logic over signature τ are sequences of symbols, where each symbol is one of the following.

1. The symbol $=$
2. An element of the infinite set $\mathcal{V} = \{x_1, x_2, x_3, \dots\}$ of *variables*
3. Constant symbols, function symbols and relation symbols in τ
4. The symbol \neg called *negation*
5. The symbol \vee called *disjunction*
6. The symbol \exists called the *existential quantifier*
7. The symbols $($ and $)$ called *parenthesis*

As always, not all such sequences are formulas; only *well formed* sequences are formulas in the logic. In order to define well formed formulas, we first need to define the set of *terms*.

Definition 3.2 The set of *terms* over signature $\tau = \{C, \mathcal{F}, \mathcal{R}\}$ is inductively defined as follows.

1. Every variable $x \in \mathcal{V}$ is a term.
2. Every constant symbol c in τ is a term.
3. If f is a k -ary function in τ and t_1, t_2, \dots, t_k are terms then $f(t_1, t_2, \dots, t_k)$ is a term.

We could capture this definition succinctly by the following BNF grammar.

$$t ::= x \mid c \mid f(t, t, \dots, t)$$

where x is a variable, c is constant symbol and f is a function symbol.

Having defined terms, we can use them to define well formed formulas (wff) or just formulas for short.

Definition 3.3 A *well formed formula* (wff) over signature τ is inductively defined as follows.

1. If t_1, t_2 are terms then $t_1 = t_2$ is a wff.
2. If t_1, t_2, \dots, t_k are terms and R is a k -ary relation symbol in τ then $R(t_1, t_2, \dots, t_k)$ is a wff.
3. If φ is a wff then $(\neg\varphi)$ is a wff.

4. If φ and ψ are wffs then $(\varphi \vee \psi)$ is a wff.
5. If φ is a wff and x is a variable then $(\exists x\varphi)$ is a wff.

More succinctly, we could capture the above definitions of terms and formulas by the following BNF grammar.

$$\varphi ::= t = t \mid R(t, t, \dots, t) \mid (\neg\varphi) \mid (\varphi \vee \varphi) \mid (\exists x\varphi)$$

where x is a variable, t is term (given by Definition 3.2, and R is a relation symbol, and x is a variable.

Atomic formulas are wffs that do not have any logical operators, i.e., either of the form $t_1 = t_2$ or $R(t_1, t_2, \dots, t_k)$, where each t_i is term and R is a k -ary relation symbol. Finally, a *literal* is formula that either atomic or the negation of an atomic formula.

It is useful to introduce logical operators in addition to those in Definition 3.3. These operators can be “syntactically” defined in terms of the operators in Definition 3.3. As in propositional logic, we can define the Boolean connectives *conjunction* as $\varphi \wedge \psi = (\neg((\neg\varphi) \vee (\neg\psi)))$, *implication* as $\varphi \rightarrow \psi = ((\neg\varphi) \vee \psi)$, *true* as $\top = (\varphi \vee (\neg\varphi))$, and *false* as $\perp = (\neg\top)$. Finally, we can define *universal quantification* as $(\forall x\varphi) = (\neg(\exists x(\neg\varphi)))$.

To avoid too many parenthesis, and at the same time have an unambiguous interpretation of formulas, we will assume the following precedence of operators (from increasing to decreasing): \neg , \wedge , \vee , \rightarrow , \forall , \exists . Thus $\forall x\forall y x = y \rightarrow \neg R(x, y)$ means $(\forall x(\forall y (x = y \rightarrow (\neg R(x, y))))$. We will also drop the outermost parentheses, and since \wedge and \vee are associative, drop parentheses in formulas involving the conjunction/disjunction of multiple formulas.

Example 3.4 Consider signature $\tau = \{R\}$ where R is a binary relation symbol. The following are formulas over this signature.

- *Reflexivity*: $\forall x R(x, x)$
- *Irreflexivity*: $\forall x (\neg R(x, x))$
- *Symmetry*: $\forall x\forall y (R(x, y) \rightarrow R(y, x))$
- *Anti-symmetry*: $\forall x\forall y ((R(x, y) \wedge R(y, x)) \rightarrow x = y)$
- *Transitivity*: $\forall x\forall y\forall z ((R(x, y) \wedge R(y, z)) \rightarrow R(x, z))$

Non-examples of formulas include $R(x)$ (R expects two arguments); x (a variable is not a formula); $(R(x, y) \vee R(z, x))$ (mismatched parentheses); $\exists x$ (x is quantified but there is no formula provided as argument).

3.2 Semantics

The semantics of formulas in any logic is defined with respect to a *model*. In the context of propositional logic, models were truth assignments to the propositions. For first order logic, models will be objects that help identify the interpretation of constants and relation symbols. Such models are called *structures*.

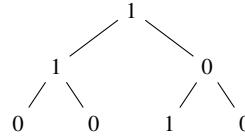


Fig. 3.1 Example of labeled binary tree.

Definition 3.5 A structure \mathcal{A} of signature τ is $\mathcal{A} = (A, \{c^{\mathcal{A}}\}_{c \in \tau}, \{f^{\mathcal{A}}\}_{f \in \tau}, \{R^{\mathcal{A}}\}_{R \in \tau})$ where

- A is a non-empty set called the *domain/universe* of the structure,
- For each constant symbol $c \in \tau$, $c^{\mathcal{A}} \in A$ is its interpretation,
- For each k -array function symbol $f \in \tau$, $f^{\mathcal{A}} : A^k \rightarrow A$ is its interpretation, and
- For each k -ary relation symbol $R \in \tau$, $R^{\mathcal{A}} \subseteq A^k$ is its interpretation.

The structure \mathcal{A} is said to be *finite* if the universe A is finite. The universe of a structure \mathcal{A} will be denoted by $u(\mathcal{A})$.

Many mathematical objects can be studied through the lens of logic. Let us look at some example signatures and structures.

Example 3.6 Consider the signature $\tau_G = \{E\}$, where E is a binary relation. We use this signature to study graphs. A graph $H = (V, E)$ modeled as a structure is $\mathcal{G} = (G, E^{\mathcal{G}})$, where the universe G is the set of vertices V , and for a pair of vertices $u, v \in G (= V)$, $E^{\mathcal{G}}uv$ ¹ holds iff $(u, v) \in E$.

Example 3.7 Let $\tau_O = \{<, S\}$ where $<$ and S are binary relation symbols. A finite order structure is $\mathcal{O} = (O, <^O, S^O)$, where O is the universe of elements, $<$ is interpreted to be an ordering relation, and S as the “successor” relation.

Example 3.8 Let $\tau_A = \{\circ\}$ where \circ is a binary function. A group would be a structure with a universe, where the operation \circ is associative, has an identity, and every element has an inverse.

Example 3.9 Labeled binary trees, where vertices are labeled by elements of Σ , can be represented as a structure in the following manner. Let $\tau_T = \{<, S_0, S_1, (Q_a)_{a \in \Sigma}\}$ where $<$, S_0 , S_1 are binary relation symbols, Q_a is a unary relation symbol. A tree (labeled by symbols in Σ) is a structure $\mathcal{T} = (T, <^T, S_0^T, S_1^T, (Q_a^T)_{a \in \Sigma})$ where elements of T are called vertices, $<$ is the ancestor relation, S_0 and S_1 are the left and right child relations, respectively, and Q_a holds in all vertices labeled by a .

For example, consider the binary tree shown in Fig. 3.1. Let us see how this tree is represented as a structure. The universe will consist of the vertices of the tree. We could use any names for the vertices. But it is convenient to name them in a manner that makes the edge relation explicit — the root will be ε , and for a vertex

¹ For a relation symbol R , we will sometimes write $R^{\mathcal{A}}a_1a_2 \dots a_n$ instead of $(a_1, a_2, \dots, a_n) \in R^{\mathcal{A}}$.

w , its left child will be $w0$, while its right child will be $w1$. Given this, the tree in Fig. 3.1 corresponds to the following structure. $\mathcal{T} = (\{\varepsilon, 0, 1, 00, 01, 10, 11\}, <^{\mathcal{T}} = \{(u, uv) \mid v \neq \varepsilon\}, S_0^{\mathcal{T}} = \{(u, u0) \mid u \in \{\varepsilon, 0, 1\}\}, S_1^{\mathcal{T}} = \{(u, u1) \mid u \in \{\varepsilon, 0, 1\}\}, Q_0 = \{1, 00, 01, 11\}, Q_1 = \{\varepsilon, 0, 10\})$.

In order to define the semantics of a first order logic formula, we need a structure, and an *assignment*. An assignment maps every variable to an element in the universe of the structure.

Definition 3.10 For a τ -structure \mathcal{A} , an *assignment* over \mathcal{A} is a function $\alpha : \mathcal{V} \rightarrow u(\mathcal{A})$ that assigns every variable $x \in \mathcal{V}$ a value $\alpha(x) \in u(\mathcal{A})$.

Fixing the values of the variable, and the interpretation of the function symbols, ensures that each term evaluates to value in $u(\mathcal{A})$. For a term t , we will abuse notation and define this value as $\alpha(t)$ and this can be defined inductively as follows.

- For a variable x , $\alpha(x)$ is simply the value α assigns to x .
- For constant symbol c , $\alpha(c) = c^{\mathcal{A}}$.
- For term $f(t_1, t_2, \dots, t_k)$, $\alpha(f(t_1, t_2, \dots, t_k)) = f^{\mathcal{A}}(\alpha(t_1), \dots, \alpha(t_k))$.

For an assignment α over \mathcal{A} , $\alpha[x \mapsto a]$ is the assignment

$$\alpha[x \mapsto a](y) = \begin{cases} \alpha(y) & \text{for } y \neq x \\ a & \text{for } x = y \end{cases}$$

We now have all the elements to define the semantics of a formula. The satisfaction relation will be a ternary relation — $\mathcal{A} \models \varphi[\alpha]$ to be read as “ φ is true/holds in \mathcal{A} under assignment α ”. The relation will be defined inductively on the structure of the formula. In defining the relation, we will also say $\mathcal{A} \not\models \varphi[\alpha]$ to mean that $\mathcal{A} \models \varphi[\alpha]$ does not hold.

Definition 3.11 The relation $\mathcal{A} \models \varphi[\alpha]$ is inductively defined as follows.

- $\mathcal{A} \models t_1 = t_2[\alpha]$ iff $\alpha(t_1) = \alpha(t_2)$
- $\mathcal{A} \models R(t_1, \dots, t_n)[\alpha]$ iff $(\alpha(t_1), \alpha(t_2), \dots, \alpha(t_n)) \in R^{\mathcal{A}}$
- $\mathcal{A} \models (\neg\varphi)[\alpha]$ iff $\mathcal{A} \not\models \varphi[\alpha]$
- $\mathcal{A} \models (\varphi \vee \psi)[\alpha]$ iff $\mathcal{A} \models \varphi[\alpha]$ or $\mathcal{A} \models \psi[\alpha]$
- $\mathcal{A} \models (\exists x\varphi)[\alpha]$ iff for some $a \in u(\mathcal{A})$, $\mathcal{A} \models \varphi[\alpha[x \mapsto a]]$

Example 3.12 Consider a structure (graph) over the vocabulary of graphs ($\tau_G = \{E\}$) $\mathcal{G} = (\{1, 2, 3, 4\}, E^{\mathcal{G}} = \{(1, 2), (2, 3), (3, 4), (4, 1)\})$. For any assignment α , $\mathcal{G} \models \forall x \exists y E(x, y)[\alpha]$ because

$$\begin{aligned} \mathcal{G} &\models \exists y E(x, y)[\alpha[x \mapsto 1]] \text{ because} \\ &\quad \mathcal{G} \models E(x, y)[\alpha[x \mapsto 1][y \mapsto 2]], \\ \mathcal{G} &\models \exists y E(x, y)[\alpha[x \mapsto 2]] \text{ because} \\ &\quad \mathcal{G} \not\models E(x, y)[\alpha[x \mapsto 2][y \mapsto 3]], \\ \mathcal{G} &\models \exists y E(x, y)[\alpha[x \mapsto 3]] \text{ because} \\ &\quad \mathcal{G} \not\models E(x, y)[\alpha[x \mapsto 3][y \mapsto 4]], \text{ and} \\ \mathcal{G} &\models \exists y E(x, y)[\alpha[x \mapsto 4]] \text{ because} \\ &\quad \mathcal{G} \not\models E(x, y)[\alpha[x \mapsto 4][y \mapsto 1]]. \end{aligned}$$

Notice that in Example 3.12, the actual assignment to variables x and y did not matter when determining the satisfaction of the formula in the graph. This is because they are *bound* by the universal and existential quantifiers in the formula φ . This leads us to the important notion of bound and free variables in a formula. We begin by defining the scope of a quantifier.

Definition 3.13 For a wff $\varphi = (\exists x\psi)$, ψ is said to be the *scope* of the quantifier $\forall x$.

Definition 3.14 Every occurrence of the variable x in $\varphi = (\exists x\psi)$ is called a *bound occurrence* of x in φ .

Any occurrence of x which is not bound is called a *free occurrence* of x in φ .

The free variables in wff φ will be denoted by $\text{free}(\varphi)$. The notation $\varphi(x_1, x_2, \dots, x_n)$ will be used to indicate that $\text{free}(\varphi) \subseteq \{x_1, \dots, x_n\}$.

Let us look at an example to understand the subtle definition of bound and free variables.

Example 3.15 Consider $\varphi = P(\mathbf{x}, \mathbf{y}) \vee (\exists x(\exists y R(x, y)) \vee Q(x, y))$ the free variables are shown in **bold**. Notice that a variable may occur both bound and free. As we will establish soon, we can change the names of bound variables without affecting the meaning of formulas. Thus $\psi = P(\mathbf{x}, \mathbf{y}) \vee (\exists u(\exists v R(u, v)) \vee Q(u, y))$ is an equivalent formula. Therefore, we will typically assume that bound and free variables are disjoint. In addition, since bound variables can be renamed without affecting its meaning, we can also assume that every bound variable is in the scope of a unique quantifier. Thus, instead of $P(\mathbf{x}, \mathbf{y}) \vee (\exists u(\exists v R(u, v)) \vee (\exists v Q(u, v)))$, we will consider the equivalent formula $P(\mathbf{x}, \mathbf{y}) \vee (\exists u(\exists v R(u, v)) \vee (\exists z Q(u, z)))$

The satisfaction of a formula in a structure \mathcal{A} under assignment α only depends on the values α assigns to the free variables; the values assigned to the bound variables in α are unimportant.

Theorem 3.16 For a formula φ and assignments α_1 and α_2 such that for every $x \in \text{free}(\varphi)$, $\alpha_1(x) = \alpha_2(x)$, $\mathcal{A} \models \varphi[\alpha_1]$ iff $\mathcal{A} \models \varphi[\alpha_2]$.

Theorem 3.16 can be proved by induction on the structure of the formula φ . The proof is left as an exercise for the reader. Theorem 3.16 suggests that if a formula has no free variables, its truth is independent of the assignment. Formulas without any free variables (i.e., those all of whose variables are bound) are an important class of formulas and have special name.

Definition 3.17 A *sentence* is a formula φ none of whose variables are free, i.e., $\text{free}(\varphi) = \emptyset$.

An immediate consequence of Theorem 3.16 is that the truth of sentences is independent of the assignment.

Proposition 3.18 For a sentence φ , and any two assignments α_1 and α_2 , $\mathcal{A} \models \varphi[\alpha_1]$ iff $\mathcal{A} \models \varphi[\alpha_2]$.

Proposition 3.18 is an immediate consequence of Theorem 3.16. Thus, for a sentence φ , we say $\mathcal{A} \models \varphi$ whenever $\mathcal{A} \models \varphi[\alpha]$ for some α .

Definition 3.19 For a sentence φ , \mathcal{A} is said to be a *model* of φ iff $\mathcal{A} \models \varphi$. We will denote by $[\![\varphi]\!]$ the set of all models of φ .

3.2.1 Satisfiability, Validity, and First order theories

Satisfiability and validity/tautologies are defined in a manner similar to that for propositional logic — a formula is satisfiable if there is some model and assignment in which it is true, and it is valid if it is true in all models and assignments.

Definition 3.20 A formula φ over signature τ is said to be *satisfiable* iff for some τ -structure \mathcal{A} and assignment α , $\mathcal{A} \models \varphi[\alpha]$.

A formula φ over signature τ is said to be logically *valid* iff for every τ -structure \mathcal{A} and assignment α , $\mathcal{A} \models \varphi[\alpha]$. We will denote this by $\models \varphi$.

We can also define with a formula φ is a *logical consequence* of a set of formulas Γ in exactly the same way as we defined it for propositional logic.

Definition 3.21 For a set of formulas Γ , we say $\mathcal{A} \models \Gamma[\alpha]$ iff for every $\varphi \in \Gamma$, $\mathcal{A} \models \varphi[\alpha]$.

We say φ is a *logical consequence* of Γ , denoted by $\Gamma \models \varphi$, if and only if for every \mathcal{A} and α , $\mathcal{A} \models \Gamma[\alpha]$ implies that $\mathcal{A} \models \varphi[\alpha]$. Thus, if $\emptyset \models \varphi$ then $\models \varphi$.

The following observation is an immediate consequence of the definition of logical consequence.

Proposition 3.22 $\Gamma \cup \{\varphi\} \models \psi$ iff $\Gamma \models \varphi \rightarrow \psi$

Finally two formulas are (semantically) equivalent, if the hold in exactly the same set of structures and assignments.

Definition 3.23 Formulas φ and ψ are said to be *logically equivalent* (denoted $\varphi \equiv \psi$) if for every \mathcal{A} and assignment α , $\mathcal{A} \models \varphi[\alpha]$ iff $\mathcal{A} \models \psi[\alpha]$.

A *first order theory* T over signature τ is any set of sentences over signature τ . A theory T is said to be *inconsistent* if there is a sentence φ such that $\{\varphi, \neg\varphi\} \subseteq T$. If T is not inconsistent then it is said to be *consistent*. Finally, T is *complete* if for every sentence φ over signature τ either $\varphi \in T$ or $\neg\varphi \in T$.

First order theories are typically identified by structures or axioms (i.e., sentences) as follows. For a structure \mathcal{A} , the first order theory of \mathcal{A} , denoted $\text{Th}(\mathcal{A})$, is defined as

$$\text{Th}(\mathcal{A}) = \{\varphi \text{ a sentence} \mid \mathcal{A} \models \varphi\}.$$

Thus, $\text{Th}(\mathcal{A})$ is the set of all sentences that are true in the structure \mathcal{A} . Notice that, since for any sentence φ exactly one of φ or $\neg\varphi$ is true in \mathcal{A} (by definition of \neg), it follows that $\text{Th}(\mathcal{A})$ is consistent and complete for any structure \mathcal{A} . For a set of

structure C , the theory of C ($\text{Th}(C)$) is set of sentences that hold in all the structures of C . That is,

$$\text{Th}(C) = \bigcap_{\mathcal{A} \in C} \text{Th}(\mathcal{A}).$$

A couple of observations about this definition are worth making. First if C is an empty set then $\text{Th}(C)$ is the set of *all* sentences and is therefore inconsistent. On the other hand, for a non-empty set C , since $\text{Th}(\mathcal{A})$ is consistent for every structure $\mathcal{A} \in C$, it follows that $\text{Th}(C)$ is also consistent. However, it may or may not be complete depending on what C .

Axioms or sets of sentences, are another way in which theories are defined. For a set of sentences Γ , the theory of Γ is given by

$$\text{Th}(\Gamma) = \{\varphi \text{ a sentence} \mid \Gamma \models \varphi\}.$$

We could define $\text{Th}(\Gamma)$ in another way. Recall that for a sentence φ , $[\![\varphi]\!]$ is the set of all structures in which φ holds. We can extend this to a set of sentences Γ by defining $[\![\Gamma]\!]$ to be the set of structures in which every sentence in Γ holds. In other words, $[\![\Gamma]\!] = \cap_{\varphi \in \Gamma} [\![\varphi]\!]$. Then $\text{Th}(\Gamma)$ is nothing but $\text{Th}([\![\Gamma]\!])$. Based on the discussion in the preceding paragraph on the consistency of the theory of a set of structures, we can conclude that $\text{Th}(\Gamma)$ is consistent if and only if $[\![\Gamma]\!]$ is non-empty. Depending on the set Γ , $\text{Th}(\Gamma)$ may or may not be complete.

3.3 Overview

There are a number of computational questions related to first order logic that we will investigate in these notes. The main ones relate to whether a sentence is true in *some* structure (*satisfiability*), in *all* structures (*validity*), and in all structures belonging to some set C over a signature. These computational questions are much harder than similar questions asked in the context of propositional logic.

Let us start with the question that is conceptually the simplest: Given a structure \mathcal{A} and sentence φ , is $\mathcal{A} \models \varphi$ or equivalently, is $\varphi \in \text{Th}(\mathcal{A})$? In the context of propositional logic the analogous question (given a truth assignment v and formula φ determine if $v \models \varphi$) is a computationally simple problem — we simply evaluate φ in v which can be done in time that is linear in the size of φ . In first order logic, it is not clear how this problem can be solved. If \mathcal{A} is a finite structure, we could simply unwind the definition of satisfaction (as we did in Example 3.12) and check if the sentence holds. When \mathcal{A} is infinite, the challenge is that existential quantifiers would require us to search in an infinite universe for a witness that the formula holds. But this begs an even more basic question, if \mathcal{A} is infinite, how is it given as input to the problem? We will consider structures \mathcal{A} that are “computable” in the sense that interpretations to constant symbols can be computed, and given representations of elements in the universe, one can compute the value of a function symbol on these arguments and one can decide if any tuple formed by these elements belongs to the

interpretation of any relation symbol in the signature. Notice that we do not require the universe $u(\mathcal{A})$ itself to be a recursive set. We will not typically worry about these computability assumptions on the structure \mathcal{A} because we will consider “standard” structures that are known to be computable in this sense. We will consider structures involving numbers and arithmetic operations, like naturals, integers, rationals, reals, equipped with the standard ordering relation and arithmetic operations of addition and multiplication.

We will begin our study of computational questions related to first order logic by investigating structures \mathcal{A} for which the set $\text{Th}(\mathcal{A})$ is decidable. Decidability results in this space are often proved by a general technique of *quantifier elimination*. A theory $\text{Th}(\mathcal{A})$ is said to *admit quantifier elimination* if for every formula φ , there is a quantifier-free formula φ' such that $\text{free}(\varphi') \subseteq \text{free}(\varphi)$ and φ' is equivalent to φ with respect to $\text{Th}(\mathcal{A})$, i.e., $\text{Th}(\mathcal{A}) \models \varphi \leftrightarrow \varphi'$ ². If the process of constructing the quantifier-free formula φ' is computable, and the problem of determining if $\psi \in \text{Th}(\mathcal{A})$ is decidable for quantifier-free formulas ψ , then composing these steps, gives a decision procedure for checking if $\varphi \in \text{Th}(\mathcal{A})$; this is often the case, and so if a theory admits quantifier elimination, then it is typically decidable. We will see that $\text{Th}((\mathbb{R}, <))$ and $\text{Th}((\mathbb{R}, 0, 1, +, <))$ admit quantifier elimination and are therefore decidable; here $<$ denotes the natural ordering on numbers and $+$ denotes addition on numbers. In fact, $\text{Th}((\mathbb{Q}, <)) = \text{Th}((\mathbb{R}, <))$ and $\text{Th}((\mathbb{Q}, 0, 1, +, <)) = \text{Th}((\mathbb{R}, 0, 1, +, <))$, and therefore these theories over the rational numbers are also decidable. The observations $\text{Th}((\mathbb{Q}, <)) = \text{Th}((\mathbb{R}, <))$ and $\text{Th}((\mathbb{Q}, 0, 1, +, <)) = \text{Th}((\mathbb{R}, 0, 1, +, <))$ demonstrate that there are non-isomorphic structures that are indistinguishable in terms of the first order sentences that they satisfy. We will see that $\text{Th}((\mathbb{N}, 0, 1, +, <))$, which is known as Presburger’s arithmetic, is also decidable. An even more surprising result is that $\text{Th}((\mathbb{R}, 0, 1, +, \times, <))$ (\times denotes multiplication on numbers) admits quantifier elimination and is decidable. This is a celebrated result due to Tarski and Seidenberg, and is beyond the scope of these notes. In contrast, both $\text{Th}((\mathbb{Q}, 0, 1, +, \times, <))$ and $\text{Th}((\mathbb{N}, 0, 1, +, \times, <))$ are not recursively enumerable. The later is a form of Gödel’s Incompleteness theorem, while the former is a result due to Robinson. Notice that even though $\text{Th}((\mathbb{Q}, 0, 1, +, <)) = \text{Th}((\mathbb{R}, 0, 1, +, <))$, $\text{Th}((\mathbb{R}, 0, 1, +, \times, <)) \neq \text{Th}((\mathbb{Q}, 0, 1, +, \times, <))$. This can be seen as follows: the sentence $\varphi = \exists x x \times x = 1 + 1$ is in $\text{Th}((\mathbb{R}, 0, 1, +, \times, <))$ (as we can take $x = \sqrt{2}$) but does not belong to $\text{Th}((\mathbb{Q}, 0, 1, +, \times, <))$.

The classical decision problem is that of determining if a sentence φ is valid. That is, given a sentence φ over signature τ , determine if φ holds in every τ -structure. This problem, on first glance, may seem very general and of little practical importance. However, this is not true. It provides a framework to study general meta-theorems in logic that are independent of a particular structure or class of structures. Equally importantly, many computational questions can be reduced to this classical problem. Suppose we want to reason about a class of structures, and the class of structures can be described by a *finite* set of sentences or *axioms* Γ . For example, suppose we want to study properties that are true about groups. Recall that groups are structures

² $\varphi \leftrightarrow \psi$ is the formula $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$.

where the universe is equipped with a binary operation $\circ : S \times S \rightarrow S$ that satisfies the following properties.

1. **Associativity:** For every a, b, c , $a \circ (b \circ c) = (a \circ b) \circ c$.
2. **Identity:** There is an element $e \in S$ such that for all a , $a \circ e = e \circ a = a$.
3. **Inverse:** For every a , there is an element a' such that $a \circ a' = a' \circ a = e$, where e is the identity.

Let Γ be the set of sentences encoding the properties of \circ being associative, and having an identity and inverses. Checking if a property φ_{uniq} that says that the identity is unique — $\forall x \forall y (\varphi_{\text{id}}(x) \wedge \varphi_{\text{id}}(y)) \rightarrow x = y$, where $\varphi_{\text{id}}(u) = \forall a a \circ u = a \wedge u \circ a = a$ — holds in every group is equivalent to checking if $\Gamma \models \varphi_{\text{uniq}}$. This question is equivalent to checking if $\models (\wedge_{\psi \in \Gamma} \psi) \rightarrow \varphi_{\text{uniq}}$, which is the classical decision problem.

One of the most important results is that the classical decision problem is recursively enumerable. This is due Gödel's completeness theorem which says that there is a sound and complete proof system (like the ones we saw for propositional logic in Chap. 2) for determining validity of first order logic sentences. Since checking if a sequence of formulas constitutes formal proof in these proof systems can be mechanized, the RE-procedure simply searches for a proof of validity. Unfortunately, the problem is RE-hard, and hence undecidable. The RE-procedure for the classical decision problem means that if Γ is an RE set of sentences then $\text{Th}(\Gamma)$ is also RE. Moreover, if $\text{Th}(\Gamma)$ is *consistent and complete* then $\text{Th}(\Gamma)$ is decidable! The decision procedure for checking if $\varphi \in \text{Th}(\Gamma)$ simply dovetails the RE-procedures for checking if $\varphi \in \text{Th}(\Gamma)$ and the procedure for checking $\neg\varphi \in \text{Th}(\Gamma)$. One of these is guaranteed to succeed since the consistency and completeness of Γ guarantees that exactly one out of φ and $\neg\varphi$ belong to $\text{Th}(\Gamma)$.

Chapter 4

Quantifier Elimination and Decidability

In this chapter, we will look at the computational problem of determining if a formula belongs to the theory of a structure. The structures we will consider involve numbers and arithmetic. We will mainly focus on structures when this problem is decidable. The key idea behind the decidability algorithms in this chapter will be *quantifier elimination*. Let us begin with this key definition. A formula φ is said to be *quantifier-free* if the quantifiers \exists and \forall do not appear in φ .

Definition 4.1 (Quantifier Elimination)

A theory Γ over signature τ admits quantifier elimination if for every formula φ over signature τ , there is a quantifier free formula φ^* such that $\text{free}(\varphi^*) \subseteq \text{free}(\varphi)$ and φ^* is equivalent to φ with respect to Γ . That is,

$$\Gamma \models \varphi \leftrightarrow \varphi^*$$

where $\varphi \leftrightarrow \psi = (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$. Another way to say this is, for every structure τ -structure \mathcal{A} such that $\mathcal{A} \models \Gamma$, and every assignment α ,

$$\mathcal{A} \models \varphi[\alpha] \text{ iff } \mathcal{A} \models \varphi^*[\alpha].$$

There are a few points worth noting about Definition 4.1. First, the free variables of quantifier-free formula φ^* is required to be a subset of the free variables of φ . Thus, in any structure \mathcal{A} that satisfies Γ , the set of assignments α that satisfy φ is exactly the same as the set of assignments that satisfy φ^* . Second, while there is no requirement that the construction of φ^* from φ be computable, it is often the case that φ^* can be effectively constructed. Hence, if in addition, for any quantifier-free sentence ψ (i.e., a Boolean combination of atomic formulas built using the constants in the signature), the problem of determining if $\Gamma \models \psi$ is decidable, then $\text{Th}(\Gamma)$ is decidable when Γ admits quantifier elimination. This is the approach we will use in this chapter to establish the decidability of $\text{Th}(\mathcal{A})$ for some structures \mathcal{A} .

Finally, observe that if Γ is *inconsistent*, then it trivially admits quantifier elimination — *any* quantifier-free formula φ^* over the same free variables as φ is (vacuously) true.

ously) equivalent to φ with respect to Γ . Let us look at a simple example of quantifier elimination.

Example 4.2 Consider the structure $(\mathbb{R}, 0, 1, +, \times, <)$. Consider the formula φ with free variables a, b, c given by $\exists x \ a \times x \times x + b \times x + c = 0$. The formula φ identifies assignments to the variables a, b, c such that the polynomial $ax^2 + bx + c$ has real roots. From high school algebra, we know that this happens when the discriminant of the polynomial is non-negative. That is, the quantifier-free formula φ^* equivalent to φ is

$$4 \times a \times c \leq b \times b$$

where by $s \leq t$ we mean the formula $(s < t) \vee (s = t)$.

We conclude this section by observing that to prove that a theory Γ admits quantifier elimination, we only need to establish this for special formulas. If every formula φ of the form $\exists x \psi$, where ψ is quantifier-free, there is a quantifier-free formula φ^* such that $\text{free}(\varphi^*) \subseteq \text{free}(\varphi)$ and φ^* is equivalent to φ with respect to Γ , then Γ admits quantifier elimination. The reason for this is that we can take any formula, express \forall quantification using \exists and negation, and starting with the innermost quantified formulas, systematically eliminate one quantifier at a time in order to eliminate all quantifiers. Hence eliminating quantifiers from formulas with a single existential quantifier is all that is needed to show that a theory admits quantifier elimination. We can make one additional simplifying assumption, if needed. We can assume that when doing quantifier elimination of $\varphi = \exists x \psi$, ψ is a conjunction of literals, i.e., a conjunction of atomic formulas or their negation. The reason is that for any arbitrary formula $\exists x \rho$, where ρ is quantifier-free, we can always treat it as a Boolean formula over atomic formulas, and write it in *disjunctive normal form* — a formula is in disjunctive normal form (DNF) if it is a disjunction of one or more conjunctions of literals — and every quantifier-free formula can be rewritten into an equivalent DNF formula. Then, we notice that \exists quantifier distributes over disjunctions, and hence we can write the formula as a disjunction of formulae of the form $\exists x \rho'$, where ρ' is a conjunction of literals. If we can do quantifier elimination on such formulae, we can simply do this for each disjunct to obtain a quantifier-free formula.

Proposition 4.3 Consider a theory Γ such that for every formula φ of the form $\exists x \bigwedge_{i=1}^n \alpha_i$, where each α_i is a literal, there is a quantifier-free formula φ^* such that $\text{free}(\varphi^*) \subseteq \text{free}(\varphi)$ and $\Gamma \models \varphi \leftrightarrow \varphi^*$. Then Γ admits quantifier elimination.

Proof We will prove this by structural induction on the formulas. In the induction below, for a formula ψ , we will denote by $\text{qe}(\psi)$ the equivalent quantifier-free formula constructed by the proof.

Base Case If ψ is an atomic formula, then simply take $\text{qe}(\psi)$ to simply be ψ itself.

Case $\psi = \neg \psi_1$ It is easy to see that ψ is equivalent to the quantifier-free formula $\neg \text{qe}(\psi_1)$.

Case $\psi = \psi_1 \vee \psi_2$ It is easy to see that ψ is equivalent to the quantifier-free formula $\text{qe}(\psi_1) \vee \text{qe}(\psi_2)$.

Case $\psi = \exists x \psi_1$ Observe that ψ is equivalent to $\exists x \text{qe}(\psi_1)$. Converting $\text{qe}(\psi_1)$ to DNF, suppose $\text{qe}(\psi_1)$ is equivalent to $\bigvee_{i=1}^k \gamma_i$, where each γ_i is a conjunction of literals. Then $\exists x \text{qe}(\psi_1)$ is equivalent to $\bigvee_{i=1}^k \exists x \gamma_i$. By our assumption, each $\exists x \gamma_i$ is equivalent to the quantifier-free formula $\text{qe}(\exists x \gamma_i)$. Thus, ψ is equivalent to the quantifier-free formula $\bigvee_{i=1}^k \text{qe}(\exists x \gamma_i)$. \square

4.1 Dense Linear Orders without Endpoints

The first structure we will look at is $(\mathbb{R}, <)$, where the universe is the set of real numbers and we have one binary relation $<$ which is interpreted as the standard ordering relation on real numbers. We will show that $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination and is decidable. The procedure to eliminate quantifiers relies on the following properties of the ordering relation $<$.

| | |
|--|---------------|
| $\forall x \neg(x < x)$ | (Irreflexive) |
| $\forall x \forall y (x < y) \rightarrow \neg(y < x)$ | (Asymmetric) |
| $\forall x \forall y \forall z ((x < y) \wedge (y < z)) \rightarrow (x < z)$ | (Transitive) |
| $\forall x \forall y (x < y) \vee (x = y) \vee (y < x)$ | (Total) |
| $\forall x \forall y (x < y) \rightarrow (\exists z (x < z) \wedge (z < y))$ | (Dense) |
| $\forall x \exists y (y < x)$ | (No Min) |
| $\forall x \exists y (x < y)$ | (No Max) |

The set of these 7 sentences will be denoted as the set DLOWE. The first 4 sentences ((Irreflexive) , (Asymmetric) , (Transitive) , and (Total)) state that $<$ is a total, strict, linear order. Equation (Dense) says that the ordering $<$ is *dense*, i.e., between any two elements one can always find a third element. The last two sentences ((No Min) and (No Max)) state that there is no minimum or maximum element.

We now show that $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination. Observe that over the signature $\{<\}$, the only atomic formulas are of the form $y < z$ or $y = z$, where y, z are variables. Our first observation shows that, in the presence of (Total) , Proposition 4.3 can be specialized even further, and we can restrict our attention to formulas without negation.

Proposition 4.4 Suppose every formula of the form $\exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is atomic of the form $x < y$, $x = y$, or $y < x$, where y is a variable that is different from x , is equivalent to a quantifier-free formula φ^* with respect to $\text{Th}((\mathbb{R}, <))$. Then $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination.

Proof By Proposition 4.3, to prove that $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination, we only need to consider formulas of the form $\varphi = \exists x \bigwedge_{i=1}^n \alpha_i$, where each α_i is a literal. We first show that (Total) allows us to eliminate negation, and so $\bigwedge_{i=1}^n \alpha_i$ is equivalent to a *positive* Boolean combination (i.e., no negations) of *atomic* formulas.

Observe that, by (Total),

$$\begin{aligned}\neg(y = z) &\equiv (y < z) \vee (z < y) \\ \neg(y < z) &\equiv (y = z) \vee (z < y)\end{aligned}$$

Thus, negation can be eliminated, and $\bigwedge_{i=1}^n \alpha_i$ is equivalent to a positive Boolean combination of atomic formulas. We can convert this into disjunctive normal form, push existential quantification inside, and see that

$$\varphi \equiv \bigvee_{i=1}^{\ell} \exists x \bigwedge_{j=1}^{k_i} \beta_{ij},$$

where each β_{ij} is an atomic formula. Thus, to prove that $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination, we can focus our attention to formulas of the form $\exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is atomic.

Consider $\psi = \exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is atomic. Observe that, by (Irreflexive), $x < x \equiv \perp$. Thus, if any $\beta_i = x < x$ then ψ is equivalent to the quantifier-free formula \perp . Next, since $x = x$ is equivalent to \top , if any $\beta_i = (x = x)$, we can drop β_i from the conjunct. Finally, if one of the β_i s is of the form $y \bowtie z$, where $\bowtie \in \{=, <\}$ and $y, z \neq x$, then we can “pull out” β_i from the quantification. This is because

$$\exists x (y \bowtie z) \wedge \rho \equiv (y \bowtie z) \wedge \exists x \rho.$$

Thus, without loss of generality, each β_i is of the form $x < y$, $x = y$ or $y < x$, for $y \neq x$ and so the proposition is established. \square

Having established Proposition 4.4, we are ready to complete the proof. Consider a formula $\varphi = \exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is either $x < y$, $x = y$, or $y < x$, for $y \neq x$. We consider two cases.

- Consider the case when there is an i and variable y such that $\beta_i = (x = y)$, i.e., one of the conjuncts is an equality constraint. Assume, without loss of generality, $\beta_1 = (x = y)$. In this case, the value for x must be the same as y . We can substitute x with y and eliminate the variable x . That is,

$$\varphi \equiv \bigwedge_{i=2}^k \beta_i [x \mapsto y]$$

- Assume that none of the conjuncts β_i are equality constraints. That is, each β_i is either $x < y$ or $y < x$ for some variable $y \neq x$. In other words, we can write φ as

$$\varphi = \exists x \left(\bigwedge_{\ell \in L} \ell < x \right) \wedge \left(\bigwedge_{u \in U} x < u \right)$$

where L and U are sets of variables. Clearly, if there is a value of x that satisfies φ with respect to an assignment α , then for every $\ell \in L$ and $u \in U$, $\alpha(\ell) < \alpha(u)$.

Conversely, when L and U are non-empty, if for an assignment α , $\alpha(\ell) < \alpha(u)$ for every $\ell \in L$ and $u \in U$, then by picking x to be a value between the the “largest” element in L and the “smallest” element in U we can show that φ holds with respect to assignment α . This can always be accomplished, since $<$ is dense. Now, if either L or U is empty, then φ can be satisfied by picking a value for x that is either very small or very large, which is possible since our structure has no minimum or maximum. This reasoning shows that there is some x satisfying the constraints if and only if every variable in L takes a value that is less than the value taken by every variable in U . Using this observation, we can say

$$\varphi \equiv \bigwedge_{\ell \in L, u \in U} \ell < u.$$

When L or U is empty, the above formula is an *empty* conjunction, which by convention is \top .

We can summarize the above observations in the main theorem for this section.

Theorem 4.5 $\text{Th}((\mathbb{R}, <))$ admits quantifier elimination.

The arguments in this section that establish Theorem 4.5, only rely on (Irreflexive), (Asymmetric), (Transitive), (Total), (Dense), (No Min), and (No Max), i.e., the sentence in DLOWE. Thus, *any* structure over the signature $\{<\}$ that satisfies all the sentence in DLOWE admits quantifier elimination. For example, since the rationals also satisfy all the properties in DLOWE, they also admit quantifier elimination. More generally, we will say structure \mathcal{A} over signature $\{<\}$ is said to be *dense linear order without endpoints* if $\mathcal{A} \models \text{DLOWE}$. Two examples of dense linear orders without endpoints are $(\mathbb{R}, <)$ and $(\mathbb{Q}, <)$. We can strengthen Theorem 4.5 as follows.

Theorem 4.6 If \mathcal{A} is a dense linear order without endpoints, then $\text{Th}(\mathcal{A})$ admits quantifier elimination.

Observe that our argument for Theorem 4.6 is *constructive*. Hence, given a formula φ over $\{<\}$, there is an algorithm that will construct the equivalent quantifier-free formula φ^* . Next, if φ is a sentence, the equivalent quantifier-free formula φ^* is also a sentence (no free variables), and therefore, just a Boolean combination of \top and \perp , which can be checked to see if it is true. Thus, for example, $\text{Th}((\mathbb{R}, <))$ and $\text{Th}((\mathbb{Q}, <))$ are decidable. It is worth observing that our procedure of constructing the quantifier-free formula, relies only on the sentences in DLOWE, and so the formula we construct is independent of the universe of the structure we are working in. Thus, $\text{Th}((\mathbb{R}, <)) = \text{Th}((\mathbb{Q}, <)) = \text{Th}(\text{DLOWE})$. This shows that there can be non-isomorphic structures (like $(\mathbb{R}, <)$ and $(\mathbb{Q}, <)$) that have the same first order theory. In general, as we shall later, for any infinite structure \mathcal{A} , it will always be the case that there are (infinitely many) different (non-isomorphic) structures \mathcal{B} that will have the same theory as \mathcal{A} . We conclude this section with the main decidability result.

Theorem 4.7 $\text{Th}(\text{DLOWE})$ is decidable. An immediate consequence of this is that $\text{Th}((\mathbb{R}, <)) = \text{Th}((\mathbb{Q}, <)) = \text{Th}(\text{DLOWE})$ are decidable.

Each quantifier eliminated by our algorithm results in a quadratic blowup (because we construct a formula that compares each variable in L with each variable in U). Thus, if a sentence of size n has m quantifiers, its equivalent quantifier-free formula has size $O(n^{2^m})$ size. This analysis does not even take into account the fact that there are steps involving the construction of a DNF formula to get removing negations, etc. Thus our procedure has a doubly exponential complexity.

4.2 Linear Arithmetic

In this section, we will extend the results of Sect. 4.1 and consider properties of numbers that involve addition along with ordering. We will look at the structures $(\mathbb{R}, 0, 1, +, <)$ and $(\mathbb{Q}, 0, 1, +, <)$, where 0 and 1 are constants representing the numbers 0 and 1, respectively, and + is the binary function symbol representing addition.

The main result of this section is captured by the following two theorems.

Theorem 4.8 *The theories $\text{Th}((\mathbb{R}, 0, 1, +, <))$ and $\text{Th}((\mathbb{Q}, 0, 1, +, <))$ admit quantifier elimination.*

Since the processes of constructing quantifier-free equivalent formulas will be effective, we will in fact get a decision procedure for these theories.

Theorem 4.9 *The theories $\text{Th}((\mathbb{R}, 0, 1, +, <))$ and $\text{Th}((\mathbb{Q}, 0, 1, +, <))$ are decidable.*

We will prove Theorem 4.8. From Proposition 4.3, to show that quantifiers can be eliminated, we only need to show that quantifiers can be eliminated from formulas φ of the form $\exists x \psi$, where ψ is a quantifier-free formula. In other words, ψ is a Boolean combination of atomic formulas. Using de Morgan's laws, we can push negations inside all the way to atomic formulas. Since $<$ on both reals and rationals satisfies (Total), we can eliminate negations like in Proposition 4.4. Recall that in our signature, atomic formulas are of the form $u \bowtie v$, where $\bowtie \in \{=, <\}$ and u and v are expressions that look like $t_1 + t_2 + \dots + t_k$ with each t_i being either a variable y , or constants 0 or 1. And,

$$\neg(u = v) \equiv (u < v) \vee (v < u) \quad \neg(u < v) \equiv (u = v) \vee (v < u).$$

Therefore, without loss of generality we may assume that ψ is a *positive* Boolean combination of atomic formulas.

Consider an atomic formula of the form $u \bowtie v$, where $\bowtie \in \{<, =\}$. We will treat these atomic formulas as equations/inequations over numbers, and use standard tricks to “solve for the variable x ”. That is, we will move all the terms involving variable x to one side with constants and other variables on the other side, and then “divide” by the coefficient of x . If x was present to begin with and does not get eliminated by this process, this will give us a formula of the form $x < u$ or $x = u$ or $x > u$, where u is a *linear expression* with rational coefficients involving the variables other than x . If x

gets eliminated or was not present to begin with, then we will get a constraint of the form $0 \bowtie u$, where $\bowtie \in \{<, =\}$ and u is a linear expression with rational coefficients involving variables other than x . Such constraints are technically not formulas in our signature, since our only constants are 0 and 1 and scalar multiplication is not a function symbol in our signature. However, this will just be an intermediate step. Before we construct the quantifier-free formula, we will get back to something that is in the legal syntax of our logic.

Let us look at an example to see what we mean by “solving for x ”.

Example 4.10 Consider the constraint $x + y + z + y + 1 < y + 1 + x + 1 + x + x$. In this solving for x , will result in the constraint

$$\frac{y}{2} + \frac{z}{2} - \frac{1}{2} < x.$$

Similarly, the constraint $x + y + x + x + z < 0$ when solved for x will result in the constraint

$$x < -\frac{y}{3} - \frac{z}{3}.$$

Based on the observations above, to show that $\text{Th}((\mathbb{R}, 0, 1, +, <))$ admits quantifier elimination, we need construct quantifier-free equivalent formula for formulas of the form $\exists x \psi$, where ψ is a positive Boolean combination of constraints of the form $x < u$, $x = u$, $u < x$, $0 = u$ or $0 < u$ where u is a linear expression with rational coefficients not mentioning x . We will present two algorithms that will eliminate quantifiers from such formulas. The first algorithm due to Fourier and Motzkin, is very similar to the approach for dense linear orders without endpoints outlined in Sect. 4.1. The second is a more efficient algorithm due Ferrante and Rackoff.

4.2.1 Fourier-Motzkin

Analogous to Proposition 4.4, we can show that we need to eliminate quantifiers only in formulas, where the quantifier-free formula in the scope of the quantifier is a conjunction of constraints involving x .

Proposition 4.11 *Suppose for every formula of the form $\exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is of the form $x < u$, $x = u$, or $u < x$, where u is a linear expression with rational coefficients involving variables other than x , is equivalent to a quantifier-free formula φ^* with respect to $\text{Th}((\mathbb{R}, 0, 1, +, <))$ (or $\text{Th}((\mathbb{Q}, 0, 1, +, <))$). Then $\text{Th}((\mathbb{R}, 0, 1, +, <))$ (or $\text{Th}((\mathbb{Q}, 0, 1, +, <))$) admits quantifier elimination.*

Proof The proof is very similar to Proposition 4.4, and we recall the main ideas, leaving the details to the reader to work out. First, based on the discussion preceding this subsection, we need to eliminate the quantifier in a formula φ of the form $\exists x \psi$, where ψ is a positive Boolean combination of constraints of the form $x < u$, $x = u$, $u < x$, $0 < u$, or $0 = u$, where x does not appear in u . We can rewrite ψ in disjunctive

normal form, push the existential quantifier inside the disjunction, and finally pull constraints of the form $0 < u$ and $0 = u$ out of the quantifier to get the result. \square

Having established Proposition 4.11, the rest of the proof is similar to Sect. 4.1. Consider a formula $\varphi = \exists x \bigwedge_{i=1}^k \beta_i$, where each β_i is either $x < u$, $x = u$, or $u < x$, for a linear expression u not involving x . We consider two cases.

- Consider the case when there is an i and expression u such that $\beta_i = (x = u)$, i.e., one of the conjuncts is an equality constraint. Assume, without loss of generality, $\beta_1 = (x = u)$. In this case, the value for x must be the same as u . We can substitute x with u and eliminate the variable x . That is,

$$\varphi \equiv \bigwedge_{i=2}^k \beta_i [x \mapsto u]$$

- Assume that none of the conjuncts β_i are equality constraints. That is, each β_i is either $x < u$ or $u < x$ for some linear expression u . In other words, we can write φ as

$$\varphi = \exists x \left(\bigwedge_{u \in L} u < x \right) \wedge \left(\bigwedge_{v \in U} x < v \right)$$

where L and U are sets of linear expressions. As in the case of dense linear orders, we can argue that φ holds if and only if, every expression in L is smaller than every expression in U . Thus,

$$\varphi \equiv \bigwedge_{u \in L, v \in U} u < v.$$

When L or U is empty, the above formula is an *empty* conjunction, which by convention is \top .

The final formula constructed by the above steps has constraints of the form $u = v$ or $u < v$, where u and v are linear expressions with rational coefficients. Such constraints are not in our signature. However, they can be rewritten into an equivalent formula in our syntax — we multiple each side by the LCM of the denominators, and rearrange terms to remove negative coefficients. We illustrate this through an example.

Example 4.12 Consider the constraint

$$\frac{y}{2} + \frac{z}{2} - \frac{1}{2} < -\frac{y}{3} - \frac{z}{3}$$

involving expressions constructed in Example 4.10. The LCM of the denominators is 6. Multiplying both sides by 6, and rearranging terms, we get the following sequence of steps.

$$\begin{aligned}
\frac{y}{2} + \frac{z}{2} - \frac{1}{2} &< -\frac{y}{3} - \frac{z}{3} \\
3y + 3z - 3 &< -2y - 2z \\
5y + 5z &< 3 \\
y + y + y + y + z + z + z + z + z &< 1 + 1 + 1
\end{aligned}$$

The last line is a formula in our syntax.

We conclude the section with an example that shows how the Fourier-Motzkin approach eliminates quantifiers.

Example 4.13 Consider the formula

$$\varphi = \forall x (0 < x) \rightarrow (1 < x + y).$$

It is easy to see that the equivalent quantifier-free formula should be $y \geq 1$, or $(1 = y) \vee (1 < y)$. Let us see how the Fourier-Motzkin method constructs this expression.

Converting the for all quantifier in terms of exists, we get $\varphi = \neg \exists x \neg (0 < x \rightarrow 1 < x + y)$. Consider $\psi = \exists x \neg (0 < x \rightarrow 1 < x + y)$. We can rewrite the implication and push the negation inside to get

$$\psi \equiv \exists x (0 < x) \wedge \neg (1 < x + y).$$

Eliminating the negation using the totality axiom, solving for x , distributing the disjunctions over the conjunction, pushing existential quantifiers in, we get

$$\begin{aligned}
\psi &\equiv \exists x (0 < x) \wedge ((1 = x + y) \vee (x + y < 1)) \\
&\equiv \exists x (0 < x) \wedge ((x = 1 - y) \vee (x < 1 - y)) \\
&\equiv \exists x [(0 < x) \wedge (x = 1 - y)] \vee [(x < 0) \wedge (x < 1 - y)] \\
&\equiv [\exists x (0 < x) \wedge (x = 1 - y)] \vee [\exists x (0 < x) \wedge (x < 1 - y)]
\end{aligned}$$

Let $\psi_1 = \exists x (0 < x) \wedge (x = 1 - y)$ and $\psi_2 = \exists x (0 < x) \wedge (x < 1 - y)$. We will eliminate the quantifier in both ψ_1 and ψ_2 to get the formula for ψ .

Since ψ_1 contains an equality constraint, we have

$$\psi_1 \equiv 0 < 1 - y \equiv y < 1.$$

In ψ_2 , we have one upper bound constraint and one lower bound constraint. So, when we eliminate the quantifier, we have

$$\psi_2 \equiv 0 < 1 - y \equiv y < 1$$

Thus,

$$\psi \equiv \psi_1 \vee \psi_2 \equiv (y < 1) \vee (y < 1) \equiv (y < 1)$$

Now, $\varphi = \neg \psi \equiv \neg (y < 1) \equiv (1 = y) \vee (1 < y)$, which is what we hoped.

4.2.2 Ferrante-Rackoff

Let us fix a formula φ of the form $\exists x \psi$, where ψ is a positive Boolean combination of constraints of the form $x < u$, $x = u$, $u < x$, $0 = u$ or $0 < u$ where u is a rational expression not mentioning x . Our goal is to eliminate the quantifier in φ . The Fourier-Motzkin algorithm relies on Proposition 4.11 which reduces the obligation to remove quantifiers to very special formulas. However, this step requires rewriting a formula to disjunctive normal form (see proof of Proposition 4.11), which can lead to an exponential blow-up. The Ferrante-Rackoff method avoids this conversion.

The key idea behind this approach is as follows. Let S be the expressions arising in constraints involving x in ψ . That is,

$$S = \{u \mid \exists \text{ constraint of the form } x = u, x < u, u < x \text{ in } \psi\}.$$

Depending on the valuation of the free variables, the expressions in S will be some rational numbers. Think of them on the “number line”. Now, x can be any number, but if two expressions u_1 and u_2 evaluate to two “consecutive” values on the number line, it doesn’t matter which value of x we pick in between u_1 and u_2 . All of them will make the atomic constraints in ψ evaluate the same way. So we can just pick $(u_1 + u_2)/2$. Now, we don’t know what *order* the expressions in S will evaluate. But we can simply instantiate x to $(u + u')/2$ for *every pair* of expressions u, u' . Since, we will also do this also for the pair u, u (in which case $(u + u)/2 = u$), this will cover x being precisely equal to one of the expressions in S . For $u, v \in S$, define

$$\psi_{\frac{u+v}{2}} = \psi[x \mapsto \frac{u+v}{2}]$$

and take

$$\psi_m = \bigvee_{u, v \in S} \psi_{\frac{u+v}{2}}.$$

To cover the range of numbers less than *all* the expressions, we can instantiate x to anything smaller than all the expressions. But instead of doing this as a substitution, we just imagine instantiating x to some large negative value, and see how the atomic formulas in ψ will evaluate. Clearly, atomic formulas of the form $x < u$ will evaluate to \top , $x = u$ will evaluate to \perp , and $u < x$ will evaluate to \perp . So we can replace the atomic formulas by these values, and get a formula $\psi_{-\infty}$. That is,

$$\psi_{-\infty} = \psi[(x < u) \mapsto \top, (x = u) \mapsto \perp, (u < x) \mapsto \perp].$$

Similarly, to cover the range of rationals larger than all the expressions, we pretend instantiating x to a value much larger than the values of all the expressions. The atomic formulas $u < x$ evaluate to \top , and the formulas of the form $x = u$ and $x < u$ evaluate to \perp . Replacing these gives the formula $\psi_{+\infty}$. That is,

$$\psi_{+\infty} = \psi[(x < u) \mapsto \perp, (x = u) \mapsto \perp, (u < x) \mapsto \top].$$

We then take the disjunction of all the above formulas to eliminate x . In other words,

$$\varphi = \exists x \psi \equiv \psi_{-\infty} \vee \psi_{+\infty} \vee \psi_m.$$

Like in the Fourier-Motzkin case, the formula as written above will not be in the syntax of our logic. But as in Example 4.12, this can be rewritten in our syntax.

Let us look at an example to see how the Ferrante-Rackoff procedure works.

Example 4.14 Consider the formula $\varphi = \forall x (0 < x) \rightarrow (1 < x + y)$, from Example 4.13. Recall that the equivalent quantifier-free formula is $y \geq 1$, or $(1 = y) \vee (1 < y)$. As in Example 4.13, we can say that $\varphi = \neg\psi$, where $\psi = \exists x \neg(0 < x \rightarrow 1 < x + y) \equiv \exists x (0 < x) \wedge ((x = 1 - y) \vee (x < 1 - y))$. Let $\rho = (0 < x) \wedge ((x = 1 - y) \vee (x < 1 - y))$ and so $\psi = \exists x \rho$.

Let us first eliminate the quantifier in ψ . Based on the rewriting of ψ , we have $S = \{0, 1 - y\}$. The expressions we need to substitute x by are “ $-\infty$ ”, “ $+\infty$ ”, 0, $1 - y$, and $\frac{1-y}{2}$. We have,

$$\begin{aligned}\rho_{-\infty} &= \perp \wedge (\perp \vee \top) \equiv \perp \\ \rho_{+\infty} &= \top \wedge (\perp \vee \perp) \equiv \perp \\ \rho_0 &= (0 < 0) \wedge ((0 = 1 - y) \vee (0 < 1 - y)) \equiv \perp \\ \rho_{1-y} &= (0 < 1 - y) \wedge ((1 - y = 1 - y) \vee (1 - y < 1 - y)) \equiv (0 < 1 - y) \equiv y < 1 \\ \rho_{\frac{1-y}{2}} &= (0 < \frac{1-y}{2}) \wedge ((\frac{1-y}{2} = 1 - y) \vee (\frac{1-y}{2} < 1 - y)) \equiv (y < 1) \wedge ((y = 1) \vee (y < 1)) \equiv (y < 1)\end{aligned}$$

Thus, we have

$$\psi \equiv \perp \vee \perp \vee \perp \vee (y < 1) \vee (y < 1) \equiv (y < 1).$$

Since $\varphi = \neg\psi$, we get $\varphi \equiv \neg(y < 1) \equiv (y = 1) \vee (1 < y)$.

Our procedures for eliminating quantifiers Sections 4.2.1 and 4.2.2, rely on properties that hold in both $(\mathbb{R}, 0, 1, +, <)$ and $(\mathbb{Q}, 0, 1, +, <)$. Further, starting with any sentence φ , the procedure (say the one by Ferrante and Rackoff) will construct the same quantifier-free formula φ^* , whether we are working with reals or rationals. Thus, $\text{Th}((\mathbb{R}, 0, 1, +, <)) = \text{Th}((\mathbb{Q}, 0, 1, +, <))$. No first order sentence can distinguish $(\mathbb{R}, 0, 1, +, <)$ and $(\mathbb{Q}, 0, 1, +, <)$.

Axiomatizations

One can ask, similar to dense linear orders without endpoints, whether there is a set of axioms/sentences that capture the property of linear arithmetic of reals/rationals. This is possible, but requires care. Chapter 3 of Calculus of Computation [?], presents such an axiomatization in parallel with the quantifier elimination procedure, and claims that they are the axioms for linear arithmetic. However, the axiomatization presented there is not complete, as it does not have any axioms for the constant 1.

In general, it is true however that for theories of a *single* structure (or more generally, for consistent and complete theories), the notions of the existence of a recursive axiomatization and decidability are synonymous. One direction is easy

— if the theory is decidable, we can simply take the theory itself as its recursive axiomatization (sounds like we are cheating, but we are not). And if there is a recursive axiomatization for a complete theory, it will follow, as we will show later (Gödel's strong completeness theorem), that the membership problem is recursively enumerable, and hence by simultaneously checking if φ or $\neg\varphi$ is in the theory, we can show the problem is decidable.

4.3 Other theories that admit quantifier elimination

There are several other important theories that admit quantifier elimination that we will not consider here.

Presburger arithmetic is $\text{PresA} = \text{Th}((\mathbb{N}, 0, 1, +, <))$, and can be shown to be decidable. However, PresA does not admit quantifier elimination. For example, one can show that there is no quantifier-free formula that is equivalent to the formula $\exists x x + x = y$, which says y is even. However, we can *extend* the signature so that it admits quantifier elimination. For each $c \in \mathbb{N}$, we will introduce a unary predicate $c|\cdot$ such that $c|x$ is true if x has a value that is a multiple of c . This extended logic does admit quantifier elimination, and leads to a decision procedure.

Another important theory that admits quantifier elimination is $\text{Th}((\mathbb{R}, 0, 1, +, \times, <))$ and is decidable. This theorem is basically due to Tarski, and is called *Tarski-Seidenberg theorem*.

Chapter 5

Lower Bounds for the Validity Problem

Church-Turing Theorem and Trakhtenbrot's Theorem

The classical decision problem or *Entscheidungsproblem* is the following: Given a sentence φ over signature τ , determine if φ is valid. The problem was popularized by David Hilbert (*das Entscheidungsproblem*, or *the decision problem*), in an attempt to lead towards the formalization of mathematics. However, what *computation* meant was not clear then. These were resolved in 1936, when Church postulated that computability is captured by a class of functions using recursion schemes, and proved that the classical decision problem was not solvable using this notion of computability. A few months later, Alan Turing, in his paper that introduced Turing machines (and started the field of theoretical computer science, or even computer science), also examined the *Entscheidungsproblem* (mentioned in the title of the paper), and showed validity of first-order logic is undecidable. Soon people realized that the notions of computing defined by Turing and Church were the same — Turing in fact showed equivalence in his paper — and the Church-Turing postulate was that the notion of computability coincided with the notion of computability defined by λ -calculus and Turing machines. The undecidability of the *Entscheidungsproblem* is credited now to both Church and Turing.

The classical decision problem, in fact, turns out to be RE -complete (and hence undecidable). We will prove the hardness of this problem in this chapter. Membership in RE will be shown later in what essentially constitutes proving Gödel's completeness theorem. In fact what we will show, which is the content of the completeness theorem, is that for any *recursive* set of sentences Γ and sentence φ , the problem of determining if $\Gamma \models \varphi$ is in RE .

In this chapter, we will also consider another problem, namely that of validity in *finite* models. That is, given a sentence φ over a signature τ , determine if for every *finite* τ -structure \mathcal{A} , $\mathcal{A} \models \varphi$. We will show that this problem of validity in finite models is coRE -complete. This result is due to Trakhtenbrot. The reason for considering this problem here is because the proof of hardness is similar to showing the hardness of the classical decision problem. coRE -hardness of validity in finite models implies that checking validity in finite models is not in RE . Consequently, there is no proof systems to establish validity in finite models! This fundamental

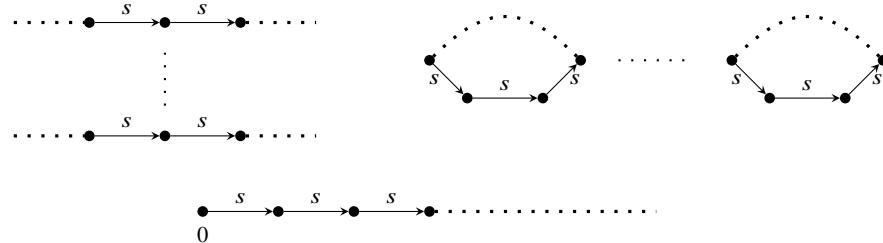


Fig. 5.1 Pictorial representation of structures satisfying (Succ) (or the sentence φ_{num}). Such structures must have a subset that is isomorphic to \mathbb{N} . They may have additional elements that form s -cycles or s -chains that are isomorphic to \mathbb{Z} (integers).

incompleteness result is easier to understand than the incompleteness result for arithmetic (Gödel's first incompleteness theorem) which we will see later.

5.1 Number Lines

Before looking at the proof of hardness, let us informally discuss some of the challenges in solving the validity problem/classical decision problem, and introduce some ideas that are central to both hardness proofs. Notice that the RE-hardness of the validity problem means that the problem of checking the *non-validity* of sentence φ , or in fact the satisfiability of $\neg\varphi$, is not recursively enumerable. This, on first reading, sometimes seems surprising. Recall that to show that φ is not valid, we need to demonstrate a structure \mathcal{A} in which φ does not hold, i.e., $\mathcal{A} \models \neg\varphi$. Can't the RE procedure for non-validity simply nondeterministically guess a structure \mathcal{A} and checking if $\mathcal{A} \models \neg\varphi$? If one can prove that φ is not valid if and only if there is a *finite* structure \mathcal{A} such that $\mathcal{A} \models \neg\varphi$, then this would indeed by a RE algorithm for non-validity. Unfortunately, it is easy to see that this not true, i.e., there are sentences ψ that satisfiable, but only in structures that are infinite. Let us look at an example.

Example 5.1 Consider the signature $\tau = \{0, s\}$, where 0 is a constant, and s is a unary function standing for *successor*. Consider the sentence

$$\varphi_{\text{succ}} = (\forall x \neg(s(x) = 0)) \wedge (\forall x \forall y (s(x) = s(y)) \rightarrow (x = y)) \quad (\text{Succ})$$

which says that s is an injective function and that the constant 0 is not the successor of any element. Consider a structure \mathcal{A} in which φ_{succ} holds. Since 0 is not the successor of any element, it follows that $0 \neq s(0)$ in \mathcal{A} . Continuing, since s is injective, $s(s(0)) \neq s(0) \neq 0$, $s(s(s(0))) \neq s(s(0)) \neq s(0) \neq 0$, and so on. We can, therefore, show by induction, that for any $i \neq j$, $s^i(0) \neq s^j(0)$ (where $s^i(0)$ is i applications of s to 0, assuming that $s^0(0) = 0$). Thus, \mathcal{A} must be infinite, because there must be a subset of its universe that is isomorphic to the natural numbers.

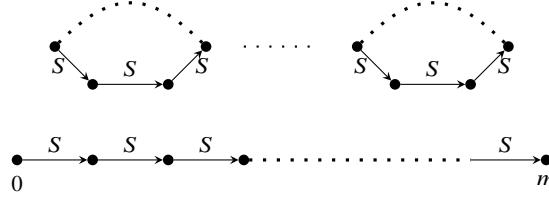


Fig. 5.2 Pictorial representation of finite structures satisfying $\varphi_{\text{fin-num}}$. Such structures must have a subset that is isomorphic to an initial segment of \mathbb{N} . They may have additional elements that form S -cycles.

Example 5.1 argues that any structure \mathcal{A} satisfying φ_{succ} (Succ) must have a subset of its universe that is isomorphic to the natural numbers. However, the universe of \mathcal{A} may have additional elements. The function s on these elements may induce sub-structures that are isomorphic to \mathbb{Z} (integers), or to cycles. A pictorial representation of such a structure is shown in Fig. 5.1, where the number of additional cycles and \mathbb{Z} -chains could be 0, finite, or infinite.

Having structure, a subset of whose universe is isomorphic to \mathbb{N} is very useful. This part of the universe can be used to model time or the number of steps of a Turing machine. It can also be used to model the indices of tape cells. If the tape symbols are encoded by numbers, then elements of this part of the universe can also be used to model tape symbols. Our reduction will use structures that have a “number line” as a sub-structure.

However, instead of using a unary successor function to create a number line, we will find it more convenient to consider structures that have a binary relation S that will represent the *graph* of the successor function s . That is, our signature will be $\{0, S\}$, where 0 is a constant as before, and S is a binary relation. We will want our binary relation S to satisfy the following sentences.

$$\begin{array}{ll}
 \forall x \exists y S(x, y) & \text{(Serial)} \\
 \forall x \forall y \forall z (S(x, y) \wedge S(x, z)) \rightarrow (y = z) & \text{(Functional)} \\
 \forall x \neg S(x, 0) & \text{(Zero)} \\
 \forall x \forall y \forall z (S(x, z) \wedge S(y, z)) \rightarrow (x = y) & \text{(Injective)}
 \end{array}$$

The first two sentences state that S is the graph of a function, the third sentence states that 0 is not the successor of any element, and the last sentence states that S is the graph of an injective function. We will denote the conjunction of these 4 sentences as φ_{num} . As observed before, structures satisfying φ_{num} will look as shown in Fig. 5.1.

Our reason for using a successor relation S , as opposed to a successor function s , to encode number lines is because this gives us the flexibility to use *partial functions* to encode an *initial segment* of \mathbb{N} . This will be used in proving the undecidability of checking validity over finite models (i.e., Trakhtenbrot’s theorem). Consider the

sentence

$$\exists m (\forall x \neg S(m, x)) \wedge (\forall x \neg(x = m) \rightarrow (\exists y S(x, y))) \quad (\text{Max})$$

which says that all elements except a maximum (m) have a successor with respect to relation S . Take $\varphi_{\text{fin-num}}$ to be the conjunction of (Max), (Functional), (Zero), and (Injective). Let \mathcal{A} be *finite* structure such that $\mathcal{A} \models \varphi_{\text{fin-num}}$. Observe that the maximum m must be an element in the successor chain starting at 0. This is because if m is not on the chain starting at 0, then since all elements except m have a S -successor, the chain starting at 0 will be infinite as argued in Example 5.1, and \mathcal{A} would not be finite. Thus any *finite* model of $\varphi_{\text{fin-num}}$ can be depicted as shown in Fig. 5.2. We will exploit this in our proof of Trakhtenbrot's theorem.

5.2 Church-Turing Theorem

In this section we will prove the following theorem.

Theorem 5.2 (Church-Turing)

Given a sentence φ over signature τ , the problem of determining if φ is valid is RE-hard.

Our proof will reduce the RE-hard language **MP** to the problem of checking validity. Recall that the universal Turing machine U recognizes the language **MP**. Without loss of generality, we will make some simplifying assumptions about U . We will assume that U has one work-tape (and an input tape); we can ignore the output tape since we are not computing a function. We assume that the input alphabet of U is $\Sigma = \{0, 1\}$ and the tape alphabet is $\Gamma = \{0, 1, \sqcup, \triangleright\}$, where \sqcup is the blank symbol, and \triangleright is the left end marker. Let Q be the set of states of U , with q_0 as the initial state, and q_{acc} as the unique accept state. The transition function δ of U is such that it ensures that input head of U never leaves the input portion of the input tape. This is ensured by moving the head to the right (+1) when the left end marker (\triangleright) is read, and moving the head left (-1) when a blank symbol (\sqcup) is read on the input tape. A configuration of U is described by current state, current input and work-tape head positions, and the contents of the work-tape.

Given binary string w , our reduction will construct a sentence φ_w such that w is accepted by U if and only if φ_w is valid. The construction will mimic the ideas in the proof of Theorem 1.23, where the sentence φ_w will describe constraints that a computation of U on w satisfies. φ_w will be sentence over the signature $\tau = \{0, S, \text{State}, \text{InpHd}, \text{TapeHd}, \text{TapeSymb}\}$, where 0 is a constant, $S, \text{State}, \text{InpHd}, \text{TapeHd}$ are binary relations, and TapeSymb is a ternary relation. The intuition behind these relation symbols in τ is as follows. 0 together with S will encode a number line as in Sect. 5.1. The remaining relation symbols are used to encode configurations of U : $\text{State}(q, t)$ holds if q is the state at time t ; $\text{InpHd}(i, t)$ and $\text{TapeHd}(j, t)$ hold if the input tape head is at cell i and work-tape head is at cell

j , at time t ; $\text{TapeSymb}(a, i, t)$ if the symbol in cell i at time t on the work-tape is a . As outlined in Sect. 5.1, we will use the structure induced by 0 and S to encode time, cell numbers, states, and tape symbols.

Like in the proof of Theorem 1.23, the sentence φ_w will state that the relations State , InpHd , TapeHd , InpSymb , TapeSymb encode the initial configuration at “time 0”, configurations at successive times follows the transition function δ , and that the accept state q_{acc} is reached at some point. In order to state these properties conveniently, we will find it convenient to make use of some auxiliary formulas that we first introduce.

- The property “Variable x stores a value which is the i th successor of 0” can be written as

$$i(x) = \exists x_1 \exists x_2 \cdots \exists x_i S(0, x_1) \wedge S(x_1, x_2) \wedge \cdots \wedge S(x_{i-1}, x_i) \wedge (x = x_i).$$

The set of states Q and the tape alphabet Γ are finite sets. We will assume that each state is q is encoded as a number in $\{0, 1, \dots, |Q| - 1\}$ and symbol a is encoded as a number in $\{0, 1, 2, 3\}$. For $b \in Q \cup \Gamma$, it would be useful to have a formula that says that “variable x stores a value which the encoding of symbol b ”. Assuming b is encoded by number i , this formula is

$$b(x) = i(x).$$

- For a formula $\psi(x, \vec{y})$ (with free variables x and \vec{y}), we will find it convenient talk about the formula when x is instantiated by $b \in Q \cup \Gamma$. We can write this as

$$\psi(b, \vec{y}) = \exists x \psi(x, \vec{y}) \wedge b(x).$$

- For a finite set S of elements (either states or tape symbols), the formula that says that a variable x takes a value in set S can be written as

$$(x \in S) = \bigvee_{b \in S} b(x).$$

- Finally, the property that “there is a unique value for x that satisfies $\psi(x, \vec{y})$ ” can be written as

$$\exists!x \psi(x, \vec{y}) = \exists x \psi(x, \vec{y}) \wedge (\forall z \psi(z, \vec{y}) \rightarrow (z = x)).$$

Let us now write a few sentences that capture properties that a valid computation of U on input w must satisfy. We begin with the condition that at time 0 the relations must encode the initial configuration. That is, the work-tape the string \triangleright , the heads pointing to cell 0 (which contains \triangleright , and the state being the initial state q_0 . Thus,

$$\begin{aligned} \varphi_{\text{initial}} = & \text{State}(q_0, 0) \wedge \text{InpHd}(0, 0) \wedge \text{TapeHd}(0, 0) \\ & \wedge (\forall c \neg(c = 0) \rightarrow \text{TapeSymb}(\sqcup, c, 0)). \end{aligned}$$

We demand that the interpretation of relations is consistent with an encoding of a configuration. That is, at all times, the state, the head positions, and symbols in each cell are unique.

$$\begin{aligned}\varphi_{\text{consistent}} = \forall t \ (\exists !x \ \text{State}(x, t)) \wedge (\exists !x \ \text{InpHd}(x, t)) \wedge (\exists !x \ \text{TapeHd}(x, t)) \\ \wedge (\forall c \exists !x \ \text{TapeSymb}(x, c, t)).\end{aligned}$$

Next, we require that configurations at successive time steps are consistent with the transition function. This is the most complicated property to write down. Consider a transition $\delta(p, a, b) = (q, d_{\text{in}}, b', d_w)$, i.e., U when in state p , reading a on the input tape, and b on the work-tape, moves to state q , writes b' on the work-tape, and moves input head in direction d_{in} and work-tape head in direction d_w . Without loss of generality, we will assume that we extend δ so that when U reaches a halting state q , $\delta(q, \cdot, \cdot)$ is defined so that the machine stays in state q ; this is so that our relations can be defined for all times. For each such tuple $(p, a, b, q, d_{\text{in}}, b', d_w)$, we will have a sentence $\varphi_{(p, a, b, q, d_{\text{in}}, b', d_w)}$ which captures when U takes a step according to this transition. To describe this property let us introduce some notation. For a tape symbol a , let S_a denote the positions on the input tape where where symbol a is written. So $S_{\text{>}} = \{0\}$, $S_{\sqcup} = \{|w| + 1\}$ ¹, and $S_a = \{i + 1 \mid w[i] = a\}$ when $a \in \{0, 1\}$ ². Finally, given a direction $d \in \{-1, +1\}$, we write $d(c, c')$ to indicate that cells c and c' are consistent with the head moving in direction d . Thus, $d(c, c') = S(c', c)$ when $d = -1$, and $d(c, c') = S(c, c')$ when $d = +1$. Given all this, we can write $\varphi_{(p, a, b, q, d_{\text{in}}, b', d_w)}$ as follows.

$$\begin{aligned}\varphi_{(p, a, b, q, d_{\text{in}}, b', d_w)} = \forall t \forall t' \forall c_{\text{in}} \forall c'_{\text{in}} \forall c_w \forall c'_w \\ (S(t, t') \wedge d_{\text{in}}(c_{\text{in}}, c'_{\text{in}}) \wedge d_w(c_w, c'_w) \wedge \text{State}(p, t) \wedge \text{InpHd}(c_{\text{in}}, t) \\ \wedge \text{TapeHd}(c_w, t) \wedge (c_{\text{in}} \in S_a) \wedge \text{TapeSymb}(b, c_w, t)) \\ \rightarrow (\text{State}(q, t') \wedge \text{InpHd}(c'_{\text{in}}, t') \wedge \text{TapeHd}(c'_w, t') \\ \wedge \text{TapeSymb}(b', c_w, t') \\ \wedge (\forall c \forall x (\neg(c = c_w) \wedge \text{TapeSymb}(x, c, t)) \rightarrow \text{TapeSymb}(x, c, t'))\end{aligned}$$

Finally, the sentence that says that every move is consistent with U 's transition function is given by

$$\varphi_{\text{transition}} = \bigvee_{\delta(p, a, b) = (q, d_{\text{in}}, b', d_w)} \varphi_{(p, a, b, q, d_{\text{in}}, b', d_w)}$$

The last condition in our sentence φ_w is the one that says that U reaches the accepting state q_{acc} . This is easy to write as

$$\varphi_{\text{accept}} = \exists t \ \text{State}(q_{\text{acc}}, t).$$

¹ Since we are assuming that the input head of U moves left when reading \sqcup , we can assume that the head never moves beyond the first \sqcup symbol on the input tape. Thus we can take $S_{\sqcup} = \{|w| + 1\}$.

² Our indices of strings start at position 0. So, $w[0]$ is written in cell 1 as cell 0 contains > .

Putting all of this together, φ_w needs to say that if 0 and S encode a number line, the computation starts in the initial configuration and follows the transition function, then U accepts. Using all the sentences we have defined, this is

$$\varphi_w = (\varphi_{\text{num}} \wedge \varphi_{\text{initial}} \wedge \varphi_{\text{consistent}} \wedge \varphi_{\text{transition}}) \rightarrow \varphi_{\text{accept}}. \quad (5.1)$$

It is easy to see that, given w , the formula φ_w can be computed by a Turing machine. To complete the proof, we need to argue that the reduction is correct. Let us consider the easy case first. Assume that φ_w is valid. Consider structure \mathcal{A} such $u(\mathcal{A}) = \mathbb{N}$. The constant 0 is interpreted as the number 0, $S(n, n')$ holds exactly when $n' = n + 1$. Next, we interpret the relations **State**, **InpHd**, **TapeHd**, **TapeSymb** in manner that is consistent with U 's computation on string w . Observe that in such a structure \mathcal{A} , φ_{num} , φ_{initial} , $\varphi_{\text{consistent}}$, and $\varphi_{\text{transition}}$ all hold. Thus, since φ_w is valid, it must be the case that φ_{accept} also holds in this model. This means that the computation of U on w reaches q_{acc} and $w \in \text{MP}$.

Let us now assume that U accepts w . Let the accepting computation of U be

$$c_0 \mapsto c_1 \mapsto \dots \mapsto c_m.$$

Our goal is to argue that φ_w is valid. Consider a structure \mathcal{A} in which φ_{num} , φ_{initial} , $\varphi_{\text{consistent}}$, and $\varphi_{\text{transition}}$ all hold. We need to argue that φ_{accept} also holds. This becomes challenging because \mathcal{A} may have additional elements that are not part of the main number line starting with 0 (see Fig. 5.1). Unfortunately, first order logic is not expressive enough to ensure that \mathcal{A} *only* contains elements of the main number line. To understand this subtlety, let us attempt to reduce $\overline{\text{MP}}$ to validity. It is tempting to think that all we need to do is to modify the sentence and demand that **State**(q_{acc}, t) does not hold for any t . That is, consider the sentence

$$\psi_w = (\varphi_{\text{num}} \wedge \varphi_{\text{initial}} \wedge \varphi_{\text{consistent}} \wedge \varphi_{\text{transition}}) \rightarrow (\forall t \neg \text{State}(q_{\text{acc}}, t)). \quad (5.2)$$

The sentence ψ_w is not valid even if U 's computation on w is non-halting — consider a structure \mathcal{B} where U 's computation faithfully encoded using **State**, **InpHd**, **TapeHd**, and **TapeSymb** on the main number line starting from 0 but **State**(q_{acc}, t) is true for t that is not on the main number line!

Coming back, the key to showing φ_{accept} holds in structure \mathcal{A} is to argue that when U accepts w and \mathcal{A} satisfies the properties in the antecedent of φ_w , then \mathcal{A} 's interpretation of **State**, **InpHd**, **TapeHd**, and **TapeSymb** is faithful to the computation on the main number line. More precisely, let the number i denote the i th successor of 0 in \mathcal{A} . We can prove by induction, that the interpretations of **State**(\cdot, i), **InpHd**(\cdot, i), **TapeHd**(\cdot, i) and **TapeSymb**(\cdot, \cdot, i) in \mathcal{A} encode the configuration c_i (i.e., the i th configuration of U 's computation on w). We leave the proof of this fact as an exercise for the reader. Having proved this, it follows that **State**(q_{acc}, m) must hold in \mathcal{A} and therefore, so does φ_{accept} . This completes the proof of Theorem 5.2.

Discussion.

Our proof shows that validity is undecidable even if the signature has one constant, 4 relation symbols and no function symbols. We could strengthen the result to the case when the signature has *only one* relation symbol — the 4 relation symbols in our reduction are modeled as a single relation which has an additional argument whose value determines which of our relations we are talking about. We could also get rid of our constant symbol 0 — we just existential quantify to get the element representing 0. Similarly, we could encode our entire reduction if our signature had only a single function symbol. It is, therefore, hard to find reasonable restrictions on the signature that make checking validity decidable. Note that quantifier alternation does play a critical role in our reduction. One could ask if there are restricted quantifier sequences that lead to decidability. A fairly complete characterization of what is decidable and what is not can be found in the book “The Classical Decision Problem” [?].

5.3 Trakhtenbrot’s Theorem

In computer science, computational problems often involve finite objects. In the context of validity, it is therefore, natural to ask how difficult is the computational problem where given a sentence φ over structure τ , we are asked to determine if φ is true in all *finite* structures. Given that the reduction in Sect. A.2 crucially relies on the sentence φ_{num} which forces the structure to be infinite, does validity become easier when considering only finite models? Clearly, *satisfiability* problem, which is the complement of the validity problem, is easy on finite structures. To determine if a sentence φ is satisfiable in a finite structure, we can simply enumerate finite models one by one, and check whether φ holds in any of them. For a finite structure \mathcal{A} , determining if $\mathcal{A} \models \varphi$ is decidable as we can simply go through our inductive definition of satisfiability to answer this question. Surprisingly, validity on finite structures is a very hard problem.

Theorem 5.3 (Trakhtenbrot)

*Given a sentence φ over signature τ , the problem of determining if φ holds in all finite structures is **coRE-hard**.*

Thus, in some sense, Theorem 5.3 says reasoning about finite structures is hard. If we want to prove theorems, then reasoning about infinite structures makes life easier as by Gödel’s completeness theorem, when theorems are true (valid), we can build machines that can identify that they are true. But over finite structures, there is no such algorithm. There is no proof system to establish theorems that hold on finite structures.

To prove Theorem 5.3 we will reduce $\overline{\text{MP}}$ to the problem of checking validity in finite models. Recall that from Theorems A.24 and A.26, we can conclude that $\overline{\text{MP}}$ is **coRE-hard** and hence, establishing such a reduction, will prove Theorem 5.3.

Again if U is the universal Turing machine accepting \mathbf{MP} , given an input w , our reduction will construct a sentence ρ_w such that w is not accepted by U if and only if ρ_w holds in all finite structures. Notice that our construction of ρ_w must differ from φ_w in (5.1) in some fundamental ways. This is because φ_w is trivially valid in all finite structures — since φ_{num} has no finite models, φ_w is vacuously true in any finite structure!

Instead of using φ_{num} , we will use $\varphi_{\text{finite-num}}$. Recall that finite structures satisfying $\varphi_{\text{finite-num}}$ look like those shown in Fig. 5.2, and so they will have a substructure that is isomorphic to an initial segment of \mathbb{N} starting at 0. Notice that the sub-structure starting from 0 has a maximum element, which is the only element that does not have an S -successor. We will make use of this as follows. To say that U does not accept w , we will say

$$\varphi_{\text{not-accept}} = \forall t (\forall x \neg S(t, x)) \rightarrow \neg \text{State}(q_{\text{acc}}, t). \quad (5.3)$$

Notice that $\varphi_{\text{not-accept}}$ requires that the state not be q_{acc} at the unique element of the structure that does not have an S -successor. On input w , our reduction will return the following sentence.

$$\rho_w = (\varphi_{\text{finite-num}} \wedge \varphi_{\text{initial}} \wedge \varphi_{\text{consistent}} \wedge \varphi_{\text{transition}}) \rightarrow \varphi_{\text{not-accept}}. \quad (5.4)$$

The proof that this is a correct reduction, is similar to the proof used in Theorem 5.2. Suppose U accepts w by computation that has k steps. Consider finite structure \mathcal{A} whose universe is $\{0, 1, \dots, k\}$, with the constant 0 being interpreted as the number 0, and $S(i, i')$ holding iff $i' = i + 1$; here k is the unique element that does not have an S -successor. Interpret the relations State , InpHd , TapeHd , and TapeSymb such that it mimics the accepting computation of U on w . Clearly \mathcal{A} satisfies the antecedent of ρ_w and $\text{State}(q_{\text{acc}}, k)$ holds. Thus $\mathcal{A} \not\models \rho_w$ and so ρ_w is not valid. Conversely, suppose U does not accept w . Consider any finite structure \mathcal{A} that satisfies the antecedent of ρ_w . Like in the proof of Theorem 5.2, one can prove by induction that the relations State , InpHd , TapeHd , and TapeSymb faithfully encode a prefix of U 's computation on w when time t is restricted to take values on sub-structure consisting of 0 and its S -successors. Notice that since U does not accept, the states at any of the times corresponding to this chain of 0 and its successors is going to be q_{acc} . Moreover, since the unique element that does not have a S -successor is guaranteed to be a successor of 0 in any finite model satisfying $\varphi_{\text{finite-num}}$, $\varphi_{\text{not-accept}}$ holds in \mathcal{A} . This establishes the correctness of the reduction.

It is worth recalling the discussion in Sect. A.2 about the formula ψ_w in (5.2) and why that does not describe a correct reduction from univL to validity. The problem there was that we may have models where $\text{State}(q_{\text{acc}}, t)$ holds for t that are not on the primary number line. Hence ψ_w may not be valid even if U does not accept w . Now, we no longer face that problem, because $\varphi_{\text{not-accept}}$ checks the state at time (i.e., the maximum) that is guaranteed to be on the main number line in any finite model of $\varphi_{\text{finite-num}}$.

Since validity and satisfiability are undecidable, a simple corollary of Theorem 5.3 is that for sentences satisfiable in finite models, there is not computable bound on the size of the model.

Corollary 5.4 *There is no computable function f such that for any sentence φ , φ is satisfiable in a finite model if and only if φ is satisfiable in a structure whose universe is bounded by $f(|\varphi|)$.*

Proof Observe that Theorem 5.3 implies that satisfiability in finite models is undecidable. Suppose (for contradiction) there was a computable function f satisfying the properties in the statement of Corollary 5.4. Then checking if a sentence is satisfiable in a finite model would be decidable as follows. Compute $f(|\varphi|)$ and check if φ holds in all finite models of size bounded by $f(|\varphi|)$. This contradicts the undecidability of finite satisfiability and hence the corollary is true. \square

Chapter 6

Incompleteness Theorems

Number theory and correctness of programs

Incompleteness results in logic argue that there are no formal systems that can prove all theorems in certain models or classes of models. In other words, for certain models or classes of models, *not all theorems have proofs, in any proof system*. We saw the first such incompleteness result in Sect. 5.3, where we proved that there is no proof system that can be used to prove that a given sentence/theorem holds in all finite structures, or special classes of finite structures like finite graphs. We, in fact, proved this incompleteness result in an equivalent form. Instead of arguing the “incompleteness” of an arbitrary proof system for finite structures, we showed that the problem of checking if a sentence holds in all finite structures is not recursively enumerable. This is the form in which we will establish incompleteness results in this chapter as well.

In this chapter we will prove two other important incompleteness results. We will show that there is no proof system that can establish the truth of all theorems (expressed in first order logic) that pertain to natural numbers with addition and multiplication. This is the famous result due to Gödel that revolutionized mathematics when it was first discovered. Natural numbers, endowed with addition and multiplication, is one of the most ubiquitous structures in mathematics and computer science. It is used to model many things in the physical/natural world and the human created world — discrete objects and others that can be discretized by approximation. For example, people are discrete objects; it’s useful to know how many children one has, or how many people can vote in a country. Goats and cows are discrete, and important in early notions of wealth and trade. Time is not discrete, but we can discretize time into intervals, like seconds, and hence use numbers to count time. Planetary positions are not discrete, but can be discretized to arcs of degree, and hence modeled as natural numbers. With discretization, a significant aspect of the physical world can be modeled as numbers, and most *observations* of physical phenomena can be modeled using numbers. Programs are also very much discrete in nature. They deal with inputs that are sequences, which can be seen as numbers; they do operations, which can be seen as similar to operations involving numbers (arithmetic/Boolean circuits); and the program itself, is a sequence of symbols that can be seen as a number.

Our second incompleteness result pertains to programs and their verification. Consider an imperative program in your favorite programming language with *assertions*. Assertions are basically properties of states that you assert in code, and are a form of *specification* asserting that the property must be true in all states where the assertion is reached. A program with assertions is said to be *correct* if in every execution of the program, whenever an assertion is reached, the asserted property holds. The problem of *program verification* is to determine whether a given program with assertions is correct. The statement that a program P with assertions is correct, is really a theorem in mathematics. And a *proof* of such a theorem, no matter what the notion of proofs are, is a mechanically checkable sequence of statements, where it should be clear that the proof asserts in the end that the program is correct. There are several sound proof systems that prove programs correct, with *Hoare logic* being a popular one. Unfortunately, as we shall see, none of these proof systems are *complete* — there are correct programs that they will fail to prove.

Incompleteness results put mathematics in a strange place than we intuitively imagined. There may be true theorems that are not provable using any set of formal rules of proof we accept. This means that there are theorems in number theory, including open problems, that may not have proofs. Similarly, there may be theorems about graphs that are unprovable, and programs whose correctness we may fail to establish.

All incompleteness results have a similar proof outline based on some form of *diagonalization*, similar to the one found by Cantor to establish the uncountability of real numbers, and similar to the one used by Turing to show the undecidability of the halting problem. Since, as computer scientists, we already know of such results, we will use these results and reductions to establish the non-recursive enumerability or incompleteness of the problems considered in this chapter.

6.1 Gödel's (First) Incompleteness Theorem

In this section we will Gödel's first incompleteness result. That is, any sound proof system for $\mathcal{N} = (\mathbb{N}, 0, 1, +, \times, <)$ is necessarily *incomplete*, i.e., for any sound proof system, there will be sentences φ such that $\mathcal{N} \models \varphi$ but φ is not provable in the proof system. However, in this section, we prove an equivalent form of this result. Namely that $\text{Th}(\mathcal{N})$ is not recursively enumerable.

Recall that (from Theorem A.16, Example A.19 and Corollary A.21) that $\overline{\text{haltL}}$ is not recursively enumerable. Our proof of the incompleteness theorem will reduce $\overline{\text{HP}}$ to the problem checking membership in $\text{Th}(\mathcal{N})$, thereby establishing that $\text{Th}(\mathcal{N})$ is not recursively enumerable. We will follow

6.2 Incompleteness of the theory of natural numbers with additional and multiplication

We want to show that the theory of natural numbers with addition and multiplication is not recursively enumerable, i.e., $\text{Th}(\mathbb{N}, 0, 1, +, \times, =, \leq)$ is not recursively enumerable.

The intuition behind this result is that the theorem stating that a program is correct (or that Turing machines halts or does not halt) *is a first-order expressible theorem in number theory!* Consequently, there is no formal proof system such that all theorems in FO arithmetic have proofs in the system.

We want to reduce the problem of Turing machine non-halting (or halting¹) to the validity problem of sentences over the theory of natural numbers. The following proof is adapted from Dexter Kozen's book "Automata and Computability".

First, we can express several interesting properties using first-order logic using addition and multiplication:

- q is the quotient and y is the remainder when x is divided by y :

$$\text{IntDiv}(x, y, q, r) : x = qy + r \wedge r < y$$

- y divides x :

$$\text{Div}(y, x) : \exists q. \text{IntDiv}(x, y, q, 0)$$

- x is prime:

$$\text{Prime}(x) : x \geq 2 \wedge \forall y. (\text{Div}(y, x) \Rightarrow (y = 1 \vee y = x))$$

- y is a power of a particular fixed prime p , i.e., $y = p^k$ for some $k \in \mathbb{N}$:

$$\text{Power}_p(y) : \forall z. ((\text{Div}(z, y) \wedge \text{Prime}(z)) \Rightarrow z = p)$$

We will now show a reduction from the non-halting problem of a Turing machine (on an empty tape) to validity of arithmetic sentences.

Given a Turing machine M with tape alphabet Γ and states Q , let us fix the alphabet $\Pi = \Gamma \cup (Q \times \Gamma)$. Let us choose a prime p larger than Π , and let us look upon sequences over Π as p -ary representations of numbers. A sequence $a_n \dots, a_0$, where each $a_i \in [0, p - 1]$ maps to the number $\sum_{i \in [1, n]} a_i p^i$.

The *computation* of M on the empty tape can be seen as a sequence of configurations $\sigma_0, \sigma_1 \dots$, where each σ is a configuration represented a word in Π^* . When M halts, configurations are bounded by some maximum length (depending on how much space M takes on the tape. Let H denote the subset of Π that has the halting state: $H = \Gamma \times HQ$, where $HQ \subseteq Q$ are the halting states. Hence M halts iff there is a sequence of configurations that represents valid moves of M that has the halting configuration, i.e., where some element of H occurs.

¹ Note that reducing the halting problem to validity also works to show non r.e.-ness since the theory is negation-complete; a negation complete theory is either decidable or not r.e.

We now encode the halting of M as the existence of a number whose p -ary representation encodes a valid halting computation of M .

A finite computation sequence $\sigma_0, \sigma_1, \dots, \sigma_n$ will be encoded as large enough blocks so that each σ_i fits into a block. If C is a large enough length to encode each configuration, we will use $c = p^C$ to capture this number. (In general, most numbers k related to the Turing machine that we need will be captured using p^k instead of k .) This number c will eventually be quantified in the formula we reduce to.

In order to say that a configuration sequence is correct, it is sufficient to demand that successive configurations are correct. In order to demand σ, σ' , two successive configurations (encoded with sequence of the same length C) are correct, it is sufficient to check every *three-element* subsequence of σ with the corresponding three-element subsequence in σ' (since Turing machines make only local changes on the tape). The three element sequences either does not encode a state (i.e., is over Γ only), in which they must be the same, or the three element sequence in σ encodes a state in the middle, in which case the corresponding three-element sequence in σ' depicts the correct evolution according to the transitions of the Turing machine.

Let V be the set of all 6-tuples $(a_1, a_2, a_3, b_1, b_2, b_3)$ that denote valid pairs of three-tuples. V includes all:

- Every $(a_1, a_2, a_3, b_1, b_2, b_3)$ such that $a_1, a_2, a_3, b_1, b_2, b_3 \in \Gamma$
- For every transition $\delta(q, a) = (b, q', R)$, the triples $(a_1, (q, a), a_2, a_1, b, (q', a_2))$, $((q, a), a_1, a_2, b, (q', a_1), a_2)$, and $(a_1, a_2, (q, a), a_1, a_2, b)$ are in V , for every $a_1, a_2 \in \Gamma$.
- For every transition $\delta(q, a) = (b, q', L)$, the triples $(a_1, (q, a), a_2, (q', a_1), b, a_2)$, $(a_1, (q, a), a_3, (q', a_1), b, a_3)$, are in V , for every $a_1, a_3 \in \Gamma$.

Note that check inconsistencies of sequences only in tuples where the “middle” symbol in the first configuration, i.e., a_2 , encodes a state.

We will write a formula that ensures that in a number encoding sequences of configurations, for every two consecutive configuration σ and σ' , every three-element subsequence in σ and the corresponding three-element subsequence in σ' , the 6 elements are related by V . This will ensure that the entire sequence of configurations is valid.

The crucial power of arithmetic with addition and multiplication is that we can encode sequences as numbers, and also *decode* sequences into their components. Here is an important formula, which says that the character in position Y of a sequence encoded by the number v is a (where $a \in [0, p - 1]$). As we said before, we encode the position Y using the number $y = p^Y$. So the following really says that the position encoded in y of the sequence encoded by v is b (assuming y is a power of p):

$$\text{Digit}(v, y, a) = \exists u. \exists r. (v = r + ay + upy \wedge r < y \wedge a < p)$$

Intuitively, let's say v 's p -ary representation can be split into $\rho_1 \cdot a \cdot \rho_2$, where $|\rho_2| = Y$. Then clearly $v = r + ap^Y + up^{Y+1}$, for some $r < y$ (where r encodes the number corresponding to ρ_2 and u encodes the number corresponding to ρ_1). Replacing p^Y by y gives $v = r + ay + upy$, which is what the above formula demands.

The intuition for the following formulae follow along similar lines, and we let the reader work this out for themselves.

We can demand that the 3-digit sequence of v at positions encoded by y are b_1, b_2 , and b_3 , using the formula:

$$\begin{aligned} 3\text{Digit}(v, y, b_1, b_2, b_3) : \exists u. \exists r. (v = r + b_1 y + b_2 p y + b_3 p p y + u p p p y \\ \wedge r < y \wedge b_1 < p \wedge b_2 < p \wedge b_3 < p) \end{aligned}$$

Now we can demand that the three digits of v at the position encoded by y match correctly the three digits of v at the position encoded by z :

$$Match(v, y, z) : \bigvee_{(a_1, a_2, a_3, b_1, b_2, b_3) \in V} (3\text{Digit}(v, y, a_1, a_2, a_3) \wedge 3\text{Digit}(v, z, b_1, b_2, b_3))$$

We can now write a formula that says that the p -ary string that represents v encodes a valid sequence of configurations of the TM evolution. For technical reasons, we will parameterize this with c and d — c is the number that encodes the size of configurations (i.e., p raised to the power of the length of configurations) and d will encode a bound on the entire length of the sequence v . The formula checks whether all pairs of three-digit sequences precisely c apart (or rather $\log_p(c)$ apart) in v match according to the Turing machine's moves, up to d :

$$ValidMoves(v, c, d) : \forall y. (Power_p(y) \wedge y p p c < d) \Rightarrow Match(v, y, yc)$$

We can state the sequence representing v starts with the initial configuration. Let $init$ be the number encoding the symbol $(q_0, \#)$, the Turing machine reading the blank symbol. Let $blank$ denote the number encoding the blank symbol $\#$. Note that the start configuration is then $(q_0, \#) \cdot \# \cdot \# \dots \#$. The following formula forces this as the first configuration of v :

$$Start(v, c) : Digit(v, 1, init) \wedge \forall y. (Power_p(y) \wedge y > 1 \wedge y < c \Rightarrow Digit(v, y, blank))$$

We can also state that the halting configuration occurs in v before d by the formula:

$$Halt(v, d) : \exists y. \left(Power_p(y) \wedge y < d \wedge \bigvee_{b \in H} Digit(v, y, b) \right)$$

We can now ready to write a formula that says that v is a valid sequence of configurations that halts. In order to do this, we first express that the number d represents an upper bound on the length of v (we then interpret v using the p -ary representation of v , with 0's padded to the left, if necessary):

$$Length(v, d) : Power_p(d) \wedge v < d$$

Note that v along with d (where $\text{Length}(v, d)$ holds) represents the precise p -ary sequence we wish to express properties about. We can now write that this sequence represents a valid halting computation:

$$\begin{aligned} \text{ValidHaltComp}(v) : \exists d. \exists c. (\text{Length}(v, d) \wedge \text{Power}_p(c) \wedge c < d \\ \wedge \text{Start}(v, c) \wedge \text{ValidMoves}(v, c, d) \wedge \text{Halt}(v, d)) \end{aligned}$$

We can now finally write a sentence that says that the Turing machine M does *not* halt:

$$\neg \exists v. \text{ValidHaltComp}(v)$$

The above formula is valid over the standard model of natural numbers with addition and multiplication iff the Turing machine does not halt. The relation $<$ can be expressed with the other relations using the following equivalence:

$$x < y \Leftrightarrow \exists z. \neg(z = 0) \wedge x + z = y$$

We hence have:

Theorem 6.1 *The first-order theory of natural numbers with addition and multiplication, $\text{Th}((\mathbb{N}, 0, 1, +, \times, =))$ is not recursively enumerable.*

6.3 Program Verification

Consider a TM that we want to check for non-halting. The TM can be realized by a program P (any programming language with infinite memory will do, as it can simulate the moves of a Turing machine; for example, a program with access to unbounded linked lists, or a program with access to an unbounded secondary storage device, or even a program that has unbounded integers). So the problem reduces to checking whether P does not halt.

Construct a program P' that is basically the program P modified so that if P halts, we add an assertion `assert false`; at the point where it halts. (An assertion of `false` doesn't hold in any program state—if you are uncomfortable with it, replace it with $x := 1$; `assert x = 0`.)

Now it is clear that the program P' satisfies its assertion if and only if P does *not* halt. Consequently, program verification is not recursively enumerable.

Theorem 6.2 *Program verification is not a recursively enumerable problem.*

6.4 Further Remarks

Gödel's proof and strengthenings

The crux of the above proof is that sequences over an alphabet, related in simple syntactic ways (like the moves of a Turing machine) can be encoded in arithmetic. Gödel was the first to discover this, and he used it in what's called Gödel numbering in order to encode *proofs* into numbers; proofs are also sequences whose validity is syntactic. This was done before a solid notion of computation (such as Turing machines or lambda calculus) existed. In fact, Gödel showed that one can encode a self-referential formula, where for any reasonable proof system, one can state in arithmetic a statement that says: "There is no proof of this statement." A proof system is damned if it proves this statement, and damned if it does not. If it proves it, then it's proved a wrong statement! And if it doesn't prove it, it's a correct statement it cannot prove! Gödel's proof combines encoding proofs/sequences as numbers and a diagonalization argument. In our proofs, we proved undecidability and hence non-r.e.-ness of Turing machine non-halting using diagonalization, and a separate proof of encoding existence of sequences into statements about numbers.

Note that the above proof extends beyond first-order arithmetic. Any logic that is more powerful than first-order arithmetic does not have a complete proof system. In fact, it turns out that the above theorem can be strengthened to show that even *quantifier-free arithmetic with addition and multiplication* (i.e., implicitly universally quantified) is undecidable and not recursively enumerable. In fact, the even simpler problem of solving Diophantine equations (given a set of polynomial equations, deciding whether there is solution using integer values) is undecidable, and checking whether there is no solution to them is not r.e.. This is a celebrated problem, called Hilbert's Tenth Problem, which was open for a long time and settled in a famous theorem by Yuri Matiyasevich in 1970.

Axiomatizations

Due to the *completeness theorem*, we know that any model whose theory is axiomatizable is *decidable*. Since the theory of arithmetic with addition and multiplication is undecidable, it follows that there can be no recursive axiomatization of it.

The Peano axioms formulated in first-order logic was an attempt to axiomatize arithmetic. It has an *infinite* set of axioms, including an axiom schema for induction, which essentially says that any first-order property about numbers (formulated as a formula with a single free variable) can be proved by induction. However, as we know from the results in this section *and* the completeness theorem, this axiom system, if sound, must be incomplete. In fact, there are natural *concrete* theorems that can be stated in FOL that are not provable in Peano arithmetic (see the results of Paris and Harrington where a version of Ramsey's theorem is shown to be unprovable in

Peano arithmetic). The *Principia Mathematica* is another formal system for which the incompleteness theorem applies, showing that there can be no recursive of it that is consistent and complete.

Returning to Program Verification: The Method using Invariants

Consider the problem of program verification again. How do people actually prove programs correct (partially correct, i.e., satisfy their assertions) in practice? The predominant method is the *invariant* method, which is basically a proof by induction. People postulate essentially a set of configurations $Inv(\bar{x})$ (called an invariant), captured as a formula in logic over a set of variables \bar{x} , and prove the following properties about it:

- The initial states are contained in Inv :

$$\forall \bar{x}. Init(\bar{x}) \Rightarrow Inv(\bar{x})$$

- If $Post(\bar{x}, \bar{x}')$ represents how the program can change configurations in a single step, then the invariant is closed under $Post$:

$$\forall \bar{x}, \bar{x}' : (Inv(\bar{x}) \wedge Post(\bar{x}, \bar{x}')) \Rightarrow Inv(\bar{x}')$$

- The invariant set and the set of *unsafe* states where the assertion is violated, do not intersect:

$$\forall \bar{x} (Inv(\bar{x}) \Rightarrow \neg Unsafe(\bar{x}))$$

Once we postulate such an invariant set, program verification boils down to proving validity of the above assertions! Clearly, if there is such an invariant set that contains all the initial states and closed under any move done by the program, then the set contains all the reachable states, and since it doesn't intersect the unsafe sets, it satisfies the assertion.

Verification methodologies such as that of Floyd and Hoare are essentially stylized proof techniques that follow the above method (expressing invariants at only loop headers or method boundaries). In fact, these stylistic methods break down (become too weak) for complex programs (such as concurrent programs or programs that pass programs/program-pointers as parameters); however, the global invariant method above is still robust and always is viable.

Now, one could ask whether such an invariant always exists. It clearly does whenever the program is correct—choose the invariant to be the set of *all* reachable states of the program: it satisfies all the above requirements.

So then where exactly lies the problem in not being able to prove a program correct? It lies in two aspects: (a) invariants may exist but may not be *expressible* in logic, and (b) invariants may be expressible in logic, but there may be no *proofs* (in a fixed formal system) that the above formulas are valid, i.e., the logic used to express the above conditions may be incomplete.

It turns out that either can happen. If we choose a weak enough logic, say a decidable logic, then we would be able to decide validity of the above formulas... but the invariant may not be expressible in logic. However, it turns out that for any reasonable programming language, the *invariant* (or the precise set of reachable states) is *always* expressible in a powerful enough logic, such as arithmetic with addition and multiplication! However, then, the logic becomes incomplete, and there may be no proofs for proving the above properties. And hence program verification, in general, remains incomplete either way. In automated program verification, one either chooses a weaker logic and builds automated decision procedures to check the properties above (this is good for shallow properties), or chooses a very expressive logic, and builds incomplete automation for logical reasoning to find proofs that establish the above properties!

Several open mathematical problems can be reduced to program verification. Consider Goldbach's conjecture, which asserts that all even integers larger than two is a sum of two primes. This problem can be reduced to program verification. Build a program that enumerates all even integers in increasing order, check if they are a sum of two primes (which can be done since both primes obviously need to be less than the considered number), and if some number isn't expressible as a sum of two primes, halt (or assert false). Then this program does not halt (or is correct) iff Goldbach's conjecture is true.

It turns out that there are similar reductions reducing the Riemann hypothesis to program verification/non-halting [?], even with a focus on building relatively small TMs.

Do Theorems have Proofs?

The incompleteness theorems raise metamathematical concerns. Is it possible that theorems may not have proofs? In any consistent formal system that is at least as expressive as arithmetic with addition and multiplication, the incompleteness theorem argues that there must be a statement (even expressible in FOL) such that either it or its negation is not provable. The methods mathematicians use is after all some formal system, which if consistent, will have unprovable theorems. For example, one could take any of the various unsolved problems in mathematics or theoretical computer science, and ask whether the problem is *independent* of the formal systems we are assuming.

Chapter 7

Quantifier-free theory of equality

In this chapter, we consider the problem of deciding the validity of sentences of the kind

$$\forall x_1, \dots, x_n. \varphi$$

where φ is quantifier free, and over an arbitrary signature.

Note that our convention for validity for formulas (with free variables) is that a formula is valid if it is true in every structure and every valuation of variables over the structure. Hence validity of the sentence $\forall x_1, \dots, x_n. \varphi$ is the same as the validity of the formula φ . Since this formula has no quantifiers, this fragment is referred to as the *quantifier-free fragment*. Furthermore, since the functions and relations are not restricted in any way, the only relation that has a fixed interpretation is *equality* (interpretation of $=$ symbol). Hence this theory is called the quantifier-free theory of equality.

We saw in the last chapter that general FOL validity is undecidable. However, the proof of that undecidability crucially used existential quantification (in particular, $\forall\exists$ quantification), and hence that proof does not apply for this fragment. We will show in fact that validity for this fragment is decidable.

Let us consider the dual problem of satisfiability. Given a quantifier-free formula $\varphi(\mathbf{x})$, is there a model and interpretation of \mathbf{x} that satisfies φ ? Note that this is the same decision problem as validity, as φ is valid iff $\neg\varphi$ is not satisfiable. Hence a decision procedure for satisfiability gives a decision procedure for validity as well.

7.1 Decidability using Bounded Models

We first make the simple observation that a quantifier-free formula φ is satisfiable iff it is satisfied in a finite model, in fact the finite model needs to be only of size n , where n is the number of *terms* mentioned in the formula.

Let $\varphi(\mathbf{x})$ be a satisfiable quantifier-free formula. Let M be a (potentially infinite) model and s be an interpretation of the variables \mathbf{x} , under which $\varphi(\mathbf{x})$ is true, i.e., $M \models \varphi(\mathbf{x})$.

Now let us construct a *finite* model M' from M that also satisfies φ . Let T be the terms mentioned in φ ; T is finite, and let's say $|T| = n$. Without loss of generality, assume $n > 0$ (if not, add a conjunction $x = x$ to the formula to introduce a term x). Let the universe in M be U_M . Now, under the interpretation s , each term $t \in T$ evaluates to an element $[t]_s \in U_M$.

Let us define the model M' as follows: the universe U' of M' is $\{[t]_s \mid t \in T\}$ (i.e., the finite subset of elements that terms evaluate to). For every relation symbol R , $R(\bar{u})$ is true in M' iff $R(\mathbf{u})$ is true in M . Functions are a bit more complex to define. Let us fix an arbitrary element e in U' . Define $f^{M'}(\bar{u})$ to be $f^M(\bar{u})$ if $f^M(\bar{u}) \in U'$, and e otherwise.

The above construction basically takes the sub-universe corresponding to the terms mentioned in φ , restricts the relations and functions to this subset, and when the function maps a vector of elements to an element outside this subset, map it to e instead.

Our claim now is that M' with the same interpretation s will also satisfy φ . First, we prove that every term in φ maps to the *same* element in M' as it does in M . And hence any atomic formula $R(\vec{t})$ as well as any atomic formula $t = t'$ will evaluate the same way under both models. It follows that the formula φ will evaluate to *true*.

The above argument shows that satisfiability for quantifier-free formulas is decidable. We can just enumerate all possible models of size n , where n is the set of terms mentioned in φ , choosing all possible interpretations for functions, relations, constants, and variables, and check if any of them satisfy φ .

If the signature (including arities) are *fixed*, and validity of formulas only over the fixed signature is to be decided, we can even do this in NP . We can just nondeterministically guess the model of size n and the interpretation, and check if the formula is satisfied. Since checking whether a formula holds in the model can be done in polynomial time, this gives an NP algorithm.

It is also easy to see that the problem is NP -hard as well, as it essentially includes Boolean logic. Reduction from SAT: Given a propositional formula α , introduce a new variable x , and replace each proposition occurrence p in φ with the atomic formula $(p=x)$. This formula is satisfiable iff α is satisfiable. Hence satisfiability of quantifier-free formulas is NP -complete.

7.2 An Algorithm for Conjunctive Formulas

The above proof that checking satisfiability is NP -hard is a bit unsatisfactory, as it shows it is hard because of the Boolean logic within in. What is the precise complexity of reasoning with equality itself?

We can ask the above question more precisely by asking what is the complexity of deciding *conjunctive* formulas. Can they be decided in polynomial time?

It turns out that the problem is indeed solvable in polynomial time. We will consider an algorithm in this section that clearly works in polynomial time if the arity of functions/relations are fixed (i.e., bounded by k). However, it turns out that one can implement this algorithm using clever data-structures to get a polynomial time algorithm even when arities are not fixed (see Calculus of Computation, Chapter 9, Section 9.3, for example).

To simplify the algorithm, let us first get rid of predicates (other than equality) from the formula. This can be done easily. Let us fix a constant π . We can model a predicate $p \subseteq U^n$ as a *function* $f_p : U^n \rightarrow U \cup \{\pi\}$, with the understanding that $p(\bar{u})$ is true iff $f_p(\bar{u}) = \pi$. Hence, given a formula φ with functions and relations, we can replace each occurrence of $p(\bar{t})$ with $f_p(\bar{t}) = \pi$, to get a formula φ' such that φ is satisfiable iff φ' is satisfiable. (We leave this as an exercise.)

Let φ be a conjunctive formula that is quantifier-free, that uses function symbols and $=$, but no relation symbols. Then $\varphi \equiv \alpha_1 \wedge \alpha_2 \wedge \dots \wedge \alpha_n$, where each α_i is a literal of the form $t = t'$ or $\neg(t = t')$.

Since the formula is conjunctive, let us look upon the formula as a set of conjuncts: $\{\alpha_1, \dots, \alpha_n\}$. In fact, let us divide these formula into two sets $E \cup D$, where E consists of equalities of the form $t = t'$ and D consists of disequalities of the form $\neg(t = t')$. Let us look upon the elements of E and D as pairs of terms of the form (t, t') .

Our key idea is now to build a model, if the formula is satisfiable, just using the terms T mentioned in the formula φ , which is finite and linear in $|\varphi|$. Furthermore, our idea is to find the *smallest* set of equality constraints that are imposed by the set of equalities in E . It turns out that such a smallest set always exists, and is called the *congruence closure* of E , denoted $CC(E)$, and is easy to compute in polynomial time. We then check whether any of the disequalities are violated in $CC(E)$. Then there is a violation iff the formula is not satisfiable.

The *congruence closure* of E , $CC(E)$ is the smallest set such that:

Includes E : For every $(t, t') \in E$, $(t, t') \in CC(E)$

Reflexive closure: For every $t \in T$, $(t, t) \in CC(E)$

Symmetric closure: For every $t, t' \in T$, if $(t, t'), (t', t) \in CC(E)$.

Transitive closure For every $t, t', t'' \in T$, if $(t, t') \in CC(E)$ and $(t', t'') \in CC(E)$, then $(t, t'') \in CC(E)$.

Congruence closure For every function symbol R of arity n , for every $f(t_1, \dots, t_n), f(t'_1, \dots, t'_n) \in T$, if $(t_1, t'_1), (t_2, t'_2), \dots, (t_n, t'_n) \in CC(E)$, then $(f(t_1, \dots, t_n), f(t'_1, \dots, t'_n)) \in CC(E)$

Intuitively, $CC(E)$ is the reflexive, symmetric and transitive closure of E , and also congruence-closed, in the sense that if two tuples of terms are deemed equal by it, then the function expressions applied on those terms are also deemed equal.

It turns out that $CC(E)$ exists (i.e., a smallest set of such equalities exists). Here is a simple proof. Define

$$\begin{aligned}
CC_0 &= E \\
CC_{i+1} &= CC_i \\
&\cup \{(t, t) \mid t \in T\} \\
&\cup \{(t, t') \mid (t', t) \in CC_i\} \\
&\cup \{(t, t') \mid \exists t'' \in T, (t, t''), (t'', t') \in CC_i\} \\
&\cup \{(f(t_1, \dots, t_n), f(t'_1, \dots, t'_n)) \mid (t_1, t'_1), \dots, (t_n, t'_n) \in CC_i, f(t_1, \dots, t_n), f(t'_1, \dots, t'_n) \in T\}
\end{aligned}$$

Now, let $CC = \cup_{i \in \mathbb{N}} CC_i$. Then it is easy to prove that CC has the required properties and is the smallest set that has these properties (readers should prove this for themselves). Since T is finite, the above procedure in fact terminates, i.e., there will be an i , such that $CC_i = CC_{i+1}$, in which case we can stop, as the future sets will all be the same. It's easy to see that this is in fact a polynomial time algorithm, since CC_i monotonically increase and can have at most $|T|^2$ pairs. We will see later in this section a concrete algorithm that does this computation a bit better.

Now let M be any model that satisfies the equalities in E . Then it is clear that M will satisfy the equalities in $CC(E)$ as well, since what we demand of $CC(E)$ are properties satisfied by equality. Consequently, the equalities in $CC(E)$ are logically implied by the equalities in E .

Let us now prove:

Lemma 7.1 *There exists a model satisfying the equalities in E and the disequalities in D exist iff $CC(E) \cap D = \emptyset$*

Proof In the forward direction, assume M is a model satisfying the equalities in E and the disequalities in D . Then, as we argued above, the equalities in $CC(E)$ must be satisfied in M as well, as they are logically implied by the equalities in E . It follows that since every disequality in D is satisfied by M , $CC(E)$ and D cannot have a common pair of terms.

In the other direction, assume $CC(E)$ and D are disjoint. Let us define the equivalence relation \sim over T defined by $CC(E)$ as $t \sim t'$ iff $(t, t') \in CC(E)$. It is easy to see that \sim is indeed an equivalence relation.

Let us construct a model M , where the universe U of M is T/\sim , i.e., the universe is the set of equivalence classes of \sim . Interpret each constant symbol c occurring in φ as the equivalence class $\llbracket c \rrbracket$, and each variable x occurring in φ as the equivalence class $\llbracket x \rrbracket$. The interpretation for a function symbol f on a vector of elements e_1, \dots, e_n is defined as follows (for the definition below, fix a particular element e^* arbitrarily):

- If there are some terms $t_1 \in e_1, \dots, t_n \in e_n$ such that $f(t_1, \dots, t_n)$ is a term in T , then map $f(e_1, \dots, e_n)$ to $\llbracket f(t_1, \dots, t_n) \rrbracket$.
- Else, map $f(e_1, \dots, e_n)$ to e^* .

Note that the above model construction is well defined only because $CC(E)$ satisfies the congruence-closure condition. For example, if $t_1 \sim t_2$, then we would need $f(t_1) \sim f(t_2)$ in order for the definition of f to be well-defined.

It is now straightforward to argue, by induction on structure of terms, that for every term t , the term evaluates to the element $\llbracket t \rrbracket$ in this model. Consequently all

the equalities E are satisfied. Also, every disequality $(t, t') \in D$, we are guaranteed $[t] \neq [t']$, since $CC(E) \cap D = \emptyset$. Hence the formula holds in the model, and is hence satisfiable. \square

The above shows that in order to check whether a quantifier-free conjunctive formula φ is satisfiable, we just need to compute $CC(E)$ and check if $CC(E) \cap D = \emptyset$, where E and D are the equalities and disequalities occurring in φ .

7.2.1 Computing $CC(E)$

One simple method for computing the congruence closure, especially on paper manually, is to compute successive equivalence relation using its *equivalence* classes.

An equivalence relation \sim over T can be represented as a set of sets that form a partition of T , i.e., as $\{E_1, \dots, E_k\}$ where each $E_i \subseteq T$, the sets are all disjoint, and their union is T .

Given such a representation of \sim , let us define a *Merge* operation on them: $Merge(\sim, E_i, E_j)$, where E_i, E_j are two equivalence classes of \sim simply merges the two equivalence classes into one and results in a new equivalence relation. More precisely, if \sim is $\{E_1, \dots, E_k\}$, then $Merge(\sim, E_i, E_j)$ is \sim' whose representation is $\{E_r \mid r \neq i, r \neq j\} \cup \{E_i \cup E_j\}$.

The algorithm for computing congruence closure is then as follows. :

- Initialize \sim to $\{\{t\} \mid t \in T\}$. In other words, we start with each term in its own equivalence class.
- For every $(t, t') \in E$, merge $[t]$ and $[t']$.
- Do the following till \sim stabilizes:
 - If there are any terms $t_1, \dots, t_n, t'_1, \dots, t'_n \in T$ such that both $f(t_1, \dots, t_n)$ and $f(t'_1, \dots, t'_n)$ are in T , and if $[f(t_1, \dots, t_n)]$ is not equal to $[f(t'_1, \dots, t'_n)]$, then merge them.
 - Check whether there is any disequality $(t, t') \in D$ such that $[t] = [t']$; if there is, report formula is unsatisfiable; otherwise, report formula is satisfiable.

Let us illustrate through an example:

Example 7.2 This example is taken from the book Calculus of Computation.

We want to know whether the following formula is satisfiable:

$$f(a, b) = a \wedge f(f(a, b), b) \neq a$$

Equalities E are $\{(f(a, b), a)\}$.

Disequalities D are $\{(f(f(a, b), b), a)\}$

The set of terms T is $\{a, b, f(a, b), f(f(a, b), b)\}$.

We start with the equivalence relation:

$$\{ \{a\}, \{b\}, \{f(a, b)\}, \{f(f(a, b), b)\} \}$$

Since $(f(a, b), a)$ are in E , we merge their equivalence classes to get:

$$\{ \{a, f(a, b)\}, \{b\}, \{f(f(a, b), b)\} \}$$

Since a and $f(a, b)$ are in the same equivalence class, and since b and b are in the same equivalence class, we merge $f(a, b)$ and $f(f(a, b), b)$ to get:

$$\{ \{a, f(a, b), f(f(a, b), b)\}, \{b\} \}$$

It is easy to verify that the equivalence classes have stabilized.

We can now check whether the disequalities are all satisfied. But since $f(f(a, b), b)$ and a are in the same equivalence class, we report that the formula is unsatisfiable.

Let us now illustrate an example where the formula is satisfiable, and also illustrate the model construction involved in the Lemma above.

Example 7.3 We want to know whether the following formula is satisfiable:

$$f(a) = b \wedge f(b) = a \wedge f(f(a)) = c \wedge \neg(f(a) = a)$$

The equalities are: $E = \{(f(a), b), (f(b), a), (f(f(a)), c)\}$.

The disequalities are: $D = \{(f(a), a)\}$

The terms are $T = \{a, b, c, f(a), f(b), f(f(a))\}$.

We start with the equivalence relation:

$$\{ \{a\}, \{b\}, \{c\}, \{f(a)\}, \{f(b)\}, \{f(f(a))\} \}$$

Since $(f(a), b)$ is in E , we merge their equivalence classes to get:

$$\{ \{a\}, \{b, f(a)\}, \{c\}, \{f(b)\}, \{f(f(a))\} \}$$

Since $(f(b), a)$ is in E , we merge their equivalence classes to get:

$$\{ \{a, f(b)\}, \{b, f(a)\}, \{c\}, \{f(f(a))\} \}$$

Since $(f(f(a)), c)$ is in E , we merge their equivalence classes to get:

$$\{ \{a, f(b)\}, \{b, f(a)\}, \{c, f(f(a))\} \}$$

Now, since b and $f(a)$ are in the same equivalence class, we must merge the equivalence classes of $f(b)$ and $f(f(a))$. We get:

$$\{ \{a, f(b), c, f(f(a))\}, \{b, f(a)\} \}$$

The equivalence class has now stabilized. We note that there is only one disequality, $(f(a), a)$, and $f(a)$ and a are in different equivalence classes. So we report the formula to be satisfiable.

Let us examine now how the proof of the Lemma above constructs a model. We have two elements in our model, e_1 standing for the class $\{a, f(b), c, f(f(a))\}$ and e_2 for the class $\{b, f(a)\}$.

Since a is in e_1 and $f(a)$ is in e_2 , we interpret that $f(e_1) = e_2$. Also, since $f(a)$ is in e_2 and $f(f(a))$ is in e_1 , we interpret $f(e_2) = e_1$. The constants a and c are interpreted as e_1 (since they belong to e_1) and the constant b is interpreted as e_2 (as b belongs to e_2). This model satisfies the formula.

There are even faster ways to do congruence closure. Notice that the key operations above have to manipulate disjoint sets and support union (merge) and find (which equivalence class does an element belong to?). This turns out to be a well-studied data structure called *disjoint-set datastructure* (or *union-find datastructure* for which efficient algorithms are known. Note that the number of equivalence classes is n , where n is the number of terms in the formula, which is of course linear in the size of the formula. Now if the signature is finite and fixed, or if the signature had functions of fixed arity (the former implies the latter), then it is easy to see that the above algorithm can be implemented in polynomial time. If the maximum arity is k , then there are only n^k possible considerations of terms to consider for identifying candidates of equivalence classes to merge. If the signature consists of functions of arbitrary arity, it turns out that one can still effect a polynomial time algorithm, but we skip this here.

SMT solvers also implement fast algorithms for congruence closure. In particular, given a quantifier-free formula (not necessarily conjunctive), they look upon the formula as a Boolean formula over propositions (each proposition being an atomic formula), and call a SAT solver to find a satisfying valuation (if the SAT solver finds it unsatisfiable, then clearly the original formula is also unsatisfiable). The satisfying valuation can now be interpreted as a *conjunctive* set of equality and disequality constraints, which can then be checked for satisfiability using congruence-closure algorithms. If this conjunctive formula is unsatisfiable, it would return a *core* set of atomic formulas that already make it unsatisfiable, and the SAT engine will add that as a clause, and continue its search for valuations.

7.3 Axioms for The Theory of Equality

In this book/course, we will treat the equality symbol ($=$) as an interpreted relation throughout—it is interpreted as equality of elements in the universe. However, in this section, we are briefly going to suspend that in order to understand what properties equality really satisfies.

Let us fix a FO signature, and for clarity, let us not have $=$ as a symbol, but instead have the symbol \doteq . If \doteq was uninterpreted, what properties would we like it to satisfy?

Here are some properties (which we will call the *congruence axioms*) that equality clearly satisfies (we continue to write relations as $t \doteq t'$), instead of $\doteq(t, t')$:

Reflexivity” $\forall x. x \dot{=} x$

Symmetry: $\forall x. (x \dot{=} y \Rightarrow y \dot{=} x)$

Transitivity: $\forall x, y, z. (x \dot{=} y \wedge y \dot{=} z) \Rightarrow x \dot{=} z$

Congruence wrt relations: For any relation R of arity n ,

$$\forall x_1, \dots, x_n, y_1, \dots, y_n. \left(\left(\bigwedge_{i \in [1, n]} x_i \dot{=} y_i \right) \Rightarrow R(x_1, \dots, x_n) \Leftrightarrow R(y_1, \dots, y_n) \right)$$

Congruence wrt functions: For any function f of arity n ,

$$\forall x_1, \dots, x_n, y_1, \dots, y_n. \left(\left(\bigwedge_{i \in [1, n]} x_i \dot{=} y_i \right) \Rightarrow f(x_1, \dots, x_n) \dot{=} f(y_1, \dots, y_n) \right)$$

First, notice that the above doesn't ensure that $\dot{=}$ will be interpreted as *true* equality on the model. For example, if there were two elements e_1 and e_2 such that all functions and relations behaved identically on them and no constant was interpreted as either of them, then we could relate them with $\dot{=}$ and satisfy all the properties above. In fact, FO with $\dot{=}$ (and without $=$) will not be able to distinguish this model from one where we removed e_2 and just had e_1 . The above properties in fact only insist that $\dot{=}$ is an equivalence relation that is also a congruence with respect to the relations and functions.

However, it turns out that the above properties are sufficient to capture equality as far as satisfiability/validity of logical formulae are concerned. It doesn't matter that the properties above capture only congruence and not true equality.

Let us formalize this. Let M be a model with universe U an interpretation of the relation $\dot{=}$ that satisfies the congruence axioms given above. Then let $M/\dot{=}$ be the model where we take as the universe the equivalence classes $U/\dot{=}$, and interpret relations and functions as follows:

- For any constant c ,

$$c^{M/\dot{=}} = [c^M]$$

- For any n -ary relation R ,

$$R^{M/\dot{=}}([e_1], [e_2], \dots, [e_n]) \text{ holds iff } R^M(e_1, \dots, e_n) \text{ holds}$$

- For any n -ary function f ,

$$f^{M/\dot{=}}([e_1], [e_2], \dots, [e_n]) = [f^M(e_1, \dots, e_n)]$$

The above says that constants are interpreted in $M/\dot{=}$ as the equivalence class of their interpretation in M , and relations and functions are interpreted in $M/\dot{=}$ depending on how the relations and functions are interpreted on the elements in their equivalence classes. The reason the above is well-defined is because $\dot{=}$ satisfies the congruence axioms.

We can now show the following:

Lemma 7.4 Fix a signature S without $=$ and \doteq . Let φ be a formula over $S \cup \{=\}$. Let φ' be φ , where $=$ is replaced with \doteq , and hence is over the signature $S \cup \{\doteq\}$.

- If φ holds in a model M , where $=$ is interpreted as equality in the model, then φ' holds in M' where \doteq is interpreted as equality, and the interpretation of \doteq does satisfy the congruence axioms.
- If φ' holds in a model M , where \doteq satisfies the congruence axioms, then φ holds in M/\doteq with $=$ interpreted as equality in the model.

We leave the above as an exercise for the reader.

A consequence of the above lemma is that a formula with $=$ is satisfiable (or valid) iff the formula, with $=$ replaced by \doteq is satisfiable (or valid) over the class of models that satisfy the congruence axioms.

Hence the above congruence axioms *define* equality as far as logic goes. Logically, there is no real difference between true equality and a congruence.

Chapter 8

Completeness Theorem: FO Validity is r.e.

Gödel proved in 1929 his first famous theorem that there is a formal proof system that can prove every valid formula in FOL. As the formal proof system one can choose a variety of proof systems (Gödel showed it for one proposed by Hilbert and Ackermann). A proof system is a formal set of rules of writing down a finite sequence (called a proof) that establishes the validity of a formula/sentence. In fact, a stronger statement is proved (let's call this the strong completeness theorem): there is a formal proof system such that for any set of axioms A , the formal system (incorporating axioms A) can prove any formula/sentence that is semantically entailed by A . In other words, the system can prove any sentence φ where φ holds in *all* models that satisfy the axioms A .

The above result is remarkable. It basically shows that any theorem that can be stated in FOL has a proof. Not only that, for any class of axiomatizable structures, the class of valid FO-formulatable theorems over such structures always has a proof. For instance, take the class of groups—they can be axiomatized using a few axioms, as we saw in Chapter 1. Consequently, every first-order formulatable theorem over groups has a proof.

Given a set of formulae/sentences A , the validity problem for the theory of A is to determine whether, given a formula/sentence φ , whether every model and every interpretation under which A holds also satisfies φ .

Connection to computability

In this book/course, we won't be studying proof systems, and hence won't prove Gödel's completeness theorem. However, we will prove essentially Gödel's completeness theorem, but where we replace proofs with *computation*.

Consider a set of axioms A which is a decidable set (i.e., it is either empty or finite or infinite where a TM can check whether a given sentence is an axiom or not). Then Gödel's completeness theorem says that every logically entailed theorem has a proof. Proofs are generally *finite* objects—they are typically finite sequences over some

signature, that closely follow some set of allowed rules, incorporating the axioms when needed, in order to prove a theorem. Whatever the proof system is, it is always true that *checking* whether such a sequence encodes a correct proof is a decidable problem. Consequently, it is easy to see that validity with respect to the axioms is a problem solvable in r.e.. This is because a Turing machine can enumerate all possible finite proofs, systematically, checking if any of them prove a given theorem, and halt if it does. So Gödel's theorem can be seen as saying that the problem of checking validity wrt any recursive set of axioms is recursively enumerable.

Our goal in this chapter is hence to prove this version of completeness. For every formula/sentence, when the TM finds that the sentence is a theorem in the theory of the given axioms, the computation itself can be viewed as a proof of the theorem.

Outline of Proof

The procedure we are going to outline is not entirely the usual one found in standard textbooks, and has a more computational flavor. As we will see, it can also be automated to some extent (we can build an r.e. procedure using calls to an SMT solver that decides the quantifier-free theory of equality).

The rough outline is as follows. We fix a countable signature. We are given a countable decidable set of axioms A and we consider the problem of proving validity of a FOL formula ψ .

1. Our procedure will work through *refutation*— we will show that $\psi' = \neg\psi$ is not satisfiable in any model satisfying the axioms. In other words, we want to show there is no model satisfying $A \cup \{\psi'\}$.
2. We first show that formulas can be translated to equivalent formulas in *prenex normal form*. Then we show that we can convert the negated formula to an *equisatisfiable* formula ψ'' over an expanded signature which has only universal quantification, and is of the form

$$\psi'' : \forall \bar{x}. \varphi$$

This process is called *Skolemization*.

3. We then show *Herbrand's Theorem* for such sets of universally quantified formulas, which roughly says that if the axioms and the formula is satisfiable, then they satisfiable in a universe that is composed of only ground terms over the signature modulo a congruence.
4. The above result will show that the universally quantified formula will be *unsatisfiable* iff replacing variables with all possible terms, which gives an infinite set of formulas, is an unsatisfiable set.
5. We will then use compactness of propositional logic to argue why this instantiated infinite set is unsatisfiable iff there is a finite subset of it that is unsatisfiable.
6. The above gives our r.e. procedure: negate the formula, Skolemize the axioms and formula, and instantiate systematically the formulas by a growing set of terms and

check whether the the resulting set is unsatisfiable. Any instantiation procedure that dovetails between the axioms and term instantiation so that all axioms are instantiated for all terms eventually will do. Each level of instantiation gives a set of quantifier-free formulae in the theory of equality, which is decidable. The algorithm will halt only if it finds that there is some level where the instantiated set is unsatisfiable.

We first show Step 2: Skolemization. Then we prove Herbrand's theorem. We then will use compactness to argue unsatisfiability can be proved using only a finite set of terms. And finally give the r.e. procedure.

8.1 Prenex Normal Form

We assume that the formulae/sentences we are considering for validity/satisfiability have first been converted to prenex normal form, i.e., to the form:

$$Q_1 x_1. Q_2 x_2 \dots Q_n x_n. \varphi$$

where φ is quantifier free, and furthermore, where no variable repeats (for every $i \neq j, x_i \neq x_j$).

We refer the reader to a standard text that shows that any formula in FOL can be converted to an equivalent formula in prenex form.

8.2 Skolemization / Herbrandization

Recall that for validity, pure universal quantification was easy to handle (we showed decidability in the last chapter). Hence, for satisfiability, pure existential quantification is easy to handle.

We can in fact eliminate all existential quantification in a satisfiability problem easily.

Consider a formula of the form

$$\psi : \forall x_1, \dots, x_n. \exists y. \varphi(x_1, \dots, x_n, y)$$

where φ is an arbitrary formula (can have quantifiers). We will show that there is an equisatisfiable formula (over an expanded signature) where we essentially remove the quantified variable x .

The formula roughly says:

For every valuation of x_1, \dots, x_n , there exists a value for y such that φ holds.

Assume there is some model that satisfies the above property. Then for every sequence of values of x_1, \dots, x_n , since there is an element y in the universe such

that φ holds, we can fix a particular choice of this element y using a new *function* f . This function in the model takes a tuple of values $(v_1, \dots, v_n) \in \mathcal{U}^n$ (standing for a valuation of x_1, \dots, x_n , respectively) and maps it to an element in \mathcal{U} . Now, instead of saying that there is a value y that satisfies φ , we can instead say that choosing y to be $f(x_1, \dots, x_n)$ satisfies φ .

More precisely, we can write instead the formula:

$$\psi' : \forall x_1, \dots, x_n. \varphi(x_1, \dots, x_n, f(x_1, \dots, x_n) / y)$$

In other words, we remove the existential quantification on y , and instead replace y in φ with $f(x_1, \dots, x_n)$. Here, f is a *new* function symbol introduced specifically for this quantification of y .¹

It should be clear that the original formula ψ is satisfiable over a signature Σ iff the above formula ψ' is satisfiable over the signature $\Sigma \cup \{f\}$, where f is a function symbol not occurring in Σ . In the forward direction, if ψ is satisfiable in a model M , we construct a model M' over the expanded signature that extends M by interpreting f on an n -tuple of values to some value y that makes φ true when x_1, \dots, x_n are evaluated as the n -tuple. This extended model M' will satisfy ψ' . In the reverse direction, if there is a model M' for ψ' , we can show that the model M which is the same as M' except with the interpretation of f erased, satisfies ψ : for every valuation of x_1, \dots, x_n , if we choose y to be $f(x_1, \dots, x_n)$, then we are guaranteed to satisfy φ .

When a formula has no universal quantification preceding an existential quantifier, the above works too, except that now the function takes *no arguments*, i.e., it is a 0-ary function. A function that takes no arguments and returns an element is essentially a constant. So we can replace such an existentially quantified variable with a new constant symbol.

More precisely, we can show that for any formula $\exists x. \varphi$ over a signature Σ , the formula $\varphi[c/x]$, where c is a new constant symbol that is not in Σ , is equisatisfiable.

The following lemma captures the above:

Lemma 8.1 *For any formula $\psi : \forall x_1, \dots, x_n. \exists y. \varphi(\bar{z}, x_1, \dots, x_n, y)$ over a signature Σ , let f be a function symbol not in Σ , and let*

$$\psi' : \forall x_1, \dots, x_n. \varphi(\bar{z}, x_1, \dots, x_n, f(x_1, \dots, x_n) / y)$$

over the signature $\Sigma \cup \{f\}$, where the arity of f is n . Then ψ and ψ' are equisatisfiable.

Also, for any formula $\psi : \exists y. \varphi(\bar{z}, y)$ over a signature Σ , let c be a constant symbol not in Σ , and let

$$\psi' : \varphi(\bar{z}, c / y)$$

over the signature $\Sigma \cup \{c\}$. Then ψ and ψ' are equisatisfiable.

Example 8.2 For example, consider the formula

¹ Some readers may wonder if we are using the axiom of choice here; we are.

$$\forall x. \exists y. R(x, y)$$

which says that every element x is R -related to *some element*. The above is equisatisfiable to the formula

$$\forall x. R(x, f(x))$$

Intuitively, the function f chooses one of the (potentially several) elements x is R -related to. Such a function exists iff every x is indeed R -related to some element.

We can now apply the above procedure of eliminating existential quantifiers repeatedly to a formula in prenex rectified normal form in order construct a purely universally quantified formula that is equisatisfiable.

Let us call formulas that are purely universally quantified *universal formulas*.

Herbrandization:

The above also shows that we can take any formula ψ and convert it into an equi-valid formula of the form $\exists x_1. \dots, \exists x_n \varphi$ over an expanded signature. We can simply take $\neg \psi$, Skolemize it to derive an equi-satisfiable formula over an expanded signature, and then take its negation, to get a formula with purely existential quantification that is equi-valid to ψ . This process of getting equi-valid formulas with existential quantification only is called Herbrandization.

8.3 Herbrand's theorem

One of the the first hurdles for solving satisfiability or proving unsatisfiability is to figure out what the *universe* for a formula might be. Clearly, we need elements to represent constants as they are terms. And we need elements to represent terms formed by applying functions (any number of times) to terms. But do we need more? Can a formula/sentence talk about elements that are are not *accessible* by using functions that involve constants?

Let us define accessible elements more formally. A *ground term* is a term without variables (it is built only using functions and constants). Let M be a model. An element e in the universe of M is said to be *accessible* if there is a ground term t such that t evaluates to the element e in the model M . A model is said to be *fully accessible* if every element of it is accessible.

We can now ask whether every sentence that is satisfiable has a fully accessible model. It turns out this is not true. For example, consider a signature that has a constant 0 and a function s and the formulae that force a number line from 0:

$$\varphi_0 : \forall x. (\neg s(x) = 0) \wedge \forall x, y. (s(x) = s(y) \Rightarrow x = y)$$

For this formula, it is indeed true that it is satisfied in a fully accessible model, for example a model that contains elements that serve as interpretations for $0, s(0), s(s(0)), \dots$ only.

However, consider adding a conjunct:

$$\varphi_1 : \varphi_0 \wedge \forall x. \exists y. (f(y) = x \wedge s(y) = y)$$

This formula means that there must be elements whose f images are $0, s(0), s(s(0)), \dots$, and hence these elements must all be *distinct* as well. These have to be different from the 0-chain and must be distinct from each other as well (as their f -images are different). Note that there are no *ground terms* that access these (infinitely many) elements. For example, one model that satisfies the above constraints is:

$$U = \mathbb{N} \cup \{i' \mid i \in \mathbb{N}\}$$

$$s(i) = i + 1, \text{ for every } i \in \mathbb{N}$$

$$s(i') = i', \text{ for every } i \in \mathbb{N}$$

$$f(i') = i, \text{ for every } i \in \mathbb{N}$$

$$f(i) = i, \text{ for every } i \in \mathbb{N}$$

Note that there are no ground terms that “evaluate” to the elements i' , where $i \in \mathbb{N}$.

It turns out however that *universal formulae do indeed have the property that satisfiable sentences always have fully accessible models*. This is called *Herbrand’s theorem* which we will prove below.

In fact, in the above example, the formula φ_0 is a universal sentence and has a fully accessible model. The sentence φ_1 does not have a fully accessible model, and notice that it uses an existential quantifier. We can, as argued in the last section, Skolemize formulas to have an equisatisfiable formula that has only universal quantification. Skolemizing φ_1 gives:

$$\varphi'_1 : \varphi_0 \wedge \forall x. (f(g(x)) = x \wedge s(g(x)) = g(x))$$

Notice that the Skolemization introduces a new function g for the existentially quantified variable y removed. And notice now there is a satisfying fully accessible model. In the model above sketched, we can make $g(i) = i'$ to satisfy the formulas (g for other elements can be defined arbitrarily).

Note that having accessible models is very pleasing. The universe can be thought of as consisting *only* of ground terms in the logic! In fact, we can *name* our elements using the terms in the logic (more precisely, equivalence classes of terms will be the elements in our universe). Also, notice that if ground terms t_1, \dots, t_n access the elements e_1, \dots, e_n , respectively, in a model M , then clearly $f(t_1, \dots, t_n)$ accesses the element $f^M(e_1, \dots, e_n)$. Consequently, in the models we build, the interpretation of functions is *fixed*— the function f will map the ground terms t_1, \dots, t_n (which are in the universe as the universe consists only of ground terms) to the ground

term $f(t_1, \dots, t_n)$. So, really, the universe, and the interpretation of constants and functions will be *fixed*. The only things to figure out are the interpretation of relations, including equality which will cause the universe to be equivalence classes over ground terms.

We now prove that this is always the case—universal sentences that are satisfiable have fully accessible models. More precisely, we will define Herbrand models where elements *are* equivalence classes of ground terms (with fixed interpretations of functions), and show that satisfiable universal sentences (also called sentences in Skolem form) have Herbrand models.

Universal Formulas and Closed Submodels

The key property that universal sentences satisfy is that they are satisfied in any *submodel* of a satisfying model, as long as the submodel is closed with respect to function applications. Let φ be a universal sentence and M , with universe U , be a model satisfying it. Let $U' \subseteq U$ that satisfies the following properties:

- For every constant c , $c^M \in U'$
- For every function symbol of arity n , if $e_1, \dots, e_n \in U'$, then $f^M(e_1, \dots, e_n) \in U'$.

Then the submodel $M' = (U', I')$ define by taking U' as the universe and interpreting all constants, functions, and relation symbols on U' exactly as in M , but restricted to U' , is called a *closed submodel*. More precisely, we define the interpretation of the closed submodel with universe U' to be:

- $c^{M'} = c^M$, for every constant symbol c
- For every relation symbol R of arity n , and for every $e_1, \dots, e_n \in U'$, $R^{M'}(e_1, \dots, e_n) \text{ iff } R^M(e_1, \dots, e_n)$
- For every function symbol f of arity n , and for every $e_1, \dots, e_n \in U'$, $f^{M'}(e_1, \dots, e_n) = f^M(e_1, \dots, e_n)$

Note that the properties that U' needs to satisfy is crucial to build the submodel—we cannot build a submodel using any sub-universe of elements (the values that f takes tuples of elements in the sub-universe to must be in the sub-universe as well).

If M' is a closed submodel of M , it turns out that M' will satisfy all the universal sentences that M satisfies (the converse does not hold, of course). Note that a sentence that has an existential quantification, say of the form $\exists x.R(x)$, may hold in M but may not hold in M' (as the elements witnessing the property may be not in the sub-universe, for example). But satisfiability is preserved for submodels on universal formulas. The proof is rather simple:

Lemma 8.3 (Closed submodel property) *Let M be a model and let M' be a closed submodel of M . Let φ be a universal sentence and let $M \models \varphi$. Then $M' \models \varphi$. Furthermore, every term evaluates to the same element in M' as it does in M .*

Proof Fix a model M , a closed submodel M' , and a universal sentence $\varphi : \forall x_1, \dots, x_n. \varphi'$ where φ' is quantifier free, where $M \models \varphi$. Let e_1, \dots, e_n be the interpretation of the variables x_1, \dots, x_n in the universe of the submodel M' . Then these belong to the universe in M , and since the universe of M' is closed under function applications, and since M' inherits the interpretations of constants and functions from M , it follows that every term t involving constants and these variables evaluate to the same element in M' as they do in M . Since M' also inherits the relations from M (including equality), it follows that all atomic formulas involving these variables evaluate to the same Boolean value in M and M' . Since φ' is quantifier-free, it too will evaluate to the same value in M as in M' . Since M' satisfies φ , for this interpretation of variables, φ' will also evaluate to *true*. Hence we have shown that for all possible interpretations of the variables in the universe of the submodel, φ' evaluates to *true*, which means that $M' \models \varphi$. \square

Note in the above that we don't make the claim for *universal formulas* but just for *universal sentences*. The reader should make sure they understand why the lemma does not hold for universal formulas.

Herbrand Models and Herbrand's Theorem

Let us now define Herbrand models.

Definition 8.4 Fix a FO signature Σ . Let GT be the set of all ground terms over Σ . A *functional congruence over ground terms* \sim is an equivalence relation over ground terms such that for every $t_1, \dots, t_n, t'_1, \dots, t'_n$, where for each $1 \leq i \leq n$, $t_i \sim t'_i$, it is the case that $f(t_1, \dots, t_n) \sim f(t'_1, \dots, t'_n)$. For such a congruence \sim , we denote the equivalence class containing t with the notation $\llbracket t \rrbracket_{\sim}$.

In the notation for equivalence classes, we sometimes write $\llbracket t \rrbracket$, if \sim is clear from context.

Definition 8.5 (Herbrand model with equality) Fix a FO signature Σ with at least one constant symbol (hence the set of ground terms over Σ is non-empty). A Herbrand model (with equality) is one where:

- The universe of the model is $U = \{\llbracket t \rrbracket \mid t \in GT\}$ consists of the set of equivalence classes of ground terms of Σ with respect to some functional congruence over ground terms \sim .
- Any constant c is interpreted as $\llbracket c \rrbracket$.
- Any function symbol f of arity n is interpreted so that for every $t_1, \dots, t_n \in GT$, $f^M(\llbracket t_1 \rrbracket, \dots, \llbracket t_n \rrbracket) = \llbracket f(t_1, \dots, t_n) \rrbracket$.

The first condition says that the universe of a Herbrand model consists of just equivalence classes of terms with respect to a functional congruence \sim . The second says that the interpretation of functions is given by the names of the elements

themselves—a function f will take the equivalence classes of n terms t_1, \dots, t_n to the equivalence class of the term $f(t_1, \dots, t_n)$. This definition of f^M is well-defined since \sim is a functional congruence over terms.²

Let us make some simple observations. First, in a Herbrand model, because of the way constants and functions are interpreted, it is easy to show, by induction, that every ground term t evaluates to the equivalence class containing it, i.e., $\llbracket t \rrbracket$. It hence follows that in a Herbrand model is fully accessible—every element $\llbracket t \rrbracket$ is accessible using the term t .

In fact the converse is also true: every fully accessible model is a Herbrand model, which will be evident in the proof of Herbrand's theorem below.

Let us now prove Herbrand's theorem.³ Herbrand's theorem states that if a universal sentence has a model, it has a Herbrand model.

The intuition of the proof is quite simple. Let M be a model satisfying a universal sentence φ . Then we can simply take the sub-universe that corresponds to all accessible elements (elements accessible using terms). Clearly, this subset is closed under function applications. And hence it defines a closed submodel that satisfies φ as well. This closed submodel, having accessible elements only, is isomorphic to a Herbrand model—we can relabel every element e using the equivalence class of *all ground terms that evaluate to the element e*, in order to make it a Herbrand model.

Theorem 8.6 (Herbrand's theorem with equality) *Let φ be a universal sentence. Then φ is satisfiable iff it is satisfiable in a Herbrand model.*

Proof We prove the forward direction (the reverse direction is trivial as if φ has a Herbrand model, then it is clearly satisfiable).

Let φ be satisfiable. Let M be a model for φ , with universe U .

Let $U' = \{t^M \mid t \text{ is a ground term}\}$. Then, clearly, U' satisfies the properties for defining a closed submodel of M —it includes the interpretations of all constant symbols, and for any function symbol of arity n and any n -tuple of elements, say t_1^M, \dots, t_n^M , it clearly contains $f^M(t_1^M, \dots, t_n^M)$, as that is precisely $f(t_1, \dots, t_n)^M$.

Now let M' be the closed submodel of M defined by U' . By the previous lemma, $M' \models \varphi$.

We now prove M' , with its elements renamed, is in fact a Herbrand model. Define a congruence on ground terms as follows: $t \sim t'$ iff $t^M = t'^M$. Verify that this is indeed a congruence on ground terms (proof: verify it is an equivalence relation, and note that if $t_1, \dots, t_n, t'_1, \dots, t'_n$ are such that each for each i , $t_i \sim t'_i$, then

² If you were a student in elementary school and you knew Herbrand models, and your math teacher asked you what $2 + 3$ is, you would say it's “ $2 + 3$ ”! The value of the function $+$ applied on the terms $\langle 2, 3 \rangle$ is simply the term/element $+(2, 3)$. You may not pass your elementary school exams though!

³ Herbrand's theorem is generally proved in a signature without equality. Then one can show that purely universally quantified formulas have a model where the elements are terms, not equivalence classes of terms. Since we want to treat equality as an interpreted relation that is always in the signature, our treatment has equivalence classes of terms. One could instead take Herbrand's theorem and introduce equality as an uninterpreted relation that satisfies the congruence axioms, and get the same result too.

it follows that $t_i^M = t_i'^M$, and hence $f(t_1, \dots, t_n)^M = f(t_1', \dots, t_n')^M$, and hence $f(t_1, \dots, t_n) \sim f(t_1', \dots, t_n')$.

Let us rename every element e as the nonempty set $\llbracket t \rrbracket$, the equivalence class of t wrt \sim , where t is some ground term that evaluates to e in M (such a term must exist since every element in U' is accessible). It is easy to prove that no two elements get named by the same equivalence class. Also, every equivalence class $\llbracket t \rrbracket$ is the label of some element in U' , namely t^M . We can easily prove, by induction on terms, that that every ground term t evaluates to $\llbracket t \rrbracket$ in M' .

It immediately follows that this is a Herbrand model satisfying φ . \square

Now, notice that in the proof of Herbrand's theorem, given a model that satisfies a formula, we built the submodel *independent* of the formula. The submodel consisted of all elements accessible using any ground term in the signature. Consequently, the same model construction works in showing that if a *set of universal sentences* S has a model, then it has a Herbrand model as well.

Corollary 8.7 (Herbrand's theorem with equality for sets of formulas) *Let S be a set of universal sentences. Then S is satisfiable iff it is satisfiable in a Herbrand model.*

8.4 Some consequences of Herbrand's theorem

Before we move to completeness, let us observe some simple consequences of Herbrand's theorem. First, it follows that if the signature is countable, then a set of sentences S has a model iff it has a *countable* model. In fact, this is true for any set of formulas as well.

Theorem 8.8 (Downward Löwenheim-Skolem Theorem) *Fix a countable signature Σ . If a set of formulas S over Σ has a model then it has a countable model.*

Proof Every formula in S can be made closed (i.e., made into a sentence) and made universal by Skolemizing it (by replacing variables by new constant symbols and removing existential quantification) to result in equisatisfiable formulas. The resulting set S' and S are equisatisfiable. In fact, it is easy to see that models for S work as models for S' , and vice-versa (we can keep the universe, and the interpretation of constants, functions, and relations in the common vocabulary the same). Let S be satisfiable. Then S' is satisfiable as well, and by Herbrand's theorem, there is a Herbrand model for S' , which, by definition, is countable. This model can be converted back to a model for S (we keep the same universe, we just throw away the interpretations of the added constants and functions during Skolemization, and interpret variables using the interpretations of constants that replaced them). Hence S has a countable model. \square

The above is a surprising result. Every set of axioms A (even infinite ones) that has a model also has a countable model. Recall that there are several complete

axiomatizations of theories where the *intended* models are uncountable. For example, there is an axiomatization of reals with addition and multiplication, i.e., for the theory of $(\mathbb{R}, 0, 1, +, \cdot)$. How then do they have a countable model? The only explanation is that even for such theories, there is a countable model that is elementary equivalent (which means it satisfies the same first-order sentences) as the model of reals with addition and multiplication! This is truly bizarre, but true!

There is a generalization of the Downward Löwenheim-Skolem Theorem, called the Löwenheim-Skolem Theorem, which we will not prove in this book, that says that if a formula over a countable signature has a satisfying model that is infinite, then it has models satisfying it of cardinalities κ . In particular, there will always be an uncountable satisfying model. This result is surprising too, as there are complete axiomatizations for certain countable models, like $(\mathbb{N}, 0, 1, +)$. The result then says that this set of axioms also has uncountable satisfying models! These are often referred to as nonstandard models of arithmetic. Again, the key thing is that such models exist, they agree with the standard model on all *first-order expressible properties*.

The above results can also be seen as saying that first-order logic is not powerful enough to talk about infinite cardinalities. A set of first-order sentences can say that the model has at most k elements, for any particular $k \in \mathbb{N}$ ($\exists x_1, \dots, x_k. \forall y \bigvee_{i \in [1, k]} (y = x_i)$). However, there is no set of first-order formulae that ensure that the models that satisfy it are countable, or have any particular cardinality. First-order logic also cannot ensure that satisfying models are *finite*—this is true since validity of first-order logic over finite models is not in r.e., but validity over arbitrary models is in r.e. (as we shall see soon in this chapter).

8.5 Gödel's completeness theorem: FO Validity is recursively enumerable

Let us fix a countable signature Σ .

An instance of the validity problem is a set (finite or infinite, but recursive) A of axioms, which are FO sentences, and a sentence φ . Our goal is to show that the problem of checking validity of such instances, i.e., checking if $A \models \varphi$, is recursively enumerable.

We first negate the formula φ . $A \models \varphi$ iff $A \cup \{\neg\varphi\}$ is unsatisfiable. Hence our goal is to prove that $S = A \cup \{\neg\varphi\}$ is unsatisfiable. We convert each formula in S to prenex rectified normal form. We then Skolemize the sentences in S to obtain a set X of sentences over an expanded signature Σ' such that S and X are equisatisfiable. Note that X is itself a recursive set. Our goal is now to show that checking whether X is unsatisfiable is recursively enumerable.

Since X has only universal formulas, we know by Herbrand's theorem that to prove X is unsatisfiable iff X has no satisfying Herbrand model.

Since the signature is countable, the set of all formulas is countable, and hence either X is finite or countable. Let us fix an enumeration of X : $\varphi_1, \varphi_2, \dots$

Now any universal formula $\forall \bar{x} \psi$ is true in a Herbrand model iff it is true when the variables \bar{x} are interpreted to be elements corresponding to all possible *ground terms*, since Herbrand models have only interpretations of ground terms in their universe. Consequently, such a universal formula is true in a Herbrand model iff for every tuple of ground terms $\bar{t}, \psi[\bar{t} / \bar{x}]$ holds in the model.

Consequently, it is easy to see that X is satisfied in a Herbrand model iff $\{\psi[\bar{t}] / \bar{x} \mid \forall \bar{x} \psi \in X, t \in GT(\Sigma)\}$ is satisfied in the Herbrand model. This leads us to:

Lemma 8.9 (Term Expansion Lemma) *A set of universal formulas Γ is satisfiable iff $\Gamma^* = \{\psi[\bar{t} / \bar{x}] \mid \forall \bar{x} \psi \in \Gamma, t \in GT(\Sigma)\}$ is satisfiable.*

Proof If Γ is satisfiable, then clearly Γ^* is satisfied in any model satisfying Γ , and hence is satisfiable as well. Conversely, if Γ^* is satisfiable, then consider a Herbrand model satisfying it (which must exist since the sentences are universal, in fact quantifier-free). Clearly, in this Herbrand model, since all elements are accessible using terms, the formulas in Γ are satisfied as well, and hence Γ is satisfiable. \square

Due to the above lemma, we can now take

$$X^* = \{\psi[\bar{t} / \bar{x} \mid \forall \bar{x} \psi \in X, t \in GT(\Sigma)\}$$

and our problem now reduces to showing X^* is unsatisfiable. Note that formulas in X^* are quantifier-free. And quantifier-free formulas admit a decidable satisfiability procedure (see previous chapter). However, even if $A = \emptyset$, X^* can be infinite. Consequently, we cannot subject the X^* to a satisfiability decision procedure.

We now want to show a *compactness theorem* for such quantifier-free sets of formulas. This will allow us to prove unsatisfiability of X^* by just looking at finite subsets of it for unsatisfiability. Note that finite subsets of X^* can be conjuncted and subject to a satisfiability decision procedure, as given in the previous chapter.

Compactness Theorem for quantifier-free grounded formulas

We want to show the following lemma:

Lemma 8.10 *Let Γ be a set of quantifier-free grounded sentences. Then Γ is satisfiable iff every finite subset of Γ is satisfiable.*

Proof If Γ is satisfiable, then, of course, every finite subset of Γ is satisfiable. We hence need to show only the converse.

We will use the propositional compactness theorem to prove this lemma. Note that since every sentence in Γ is grounded, each atomic formula is of the form $t = t'$ or $R(t_1, \dots, t_n)$, where t, t', t_1, \dots, t_n are grounded terms.

Let us introduce a set of propositions p_a for every atomic grounded formula a . We can now form a set Γ_p that contains the propositional abstraction of formulas in Γ , obtained by replacing every atomic formula a in any formula in Γ with the proposition p_a .

Now, of course, an arbitrary satisfying assignment of Γ_p may not correspond to a way of satisfying Γ , since equalities obey a set of properties, namely the *congruence axioms* detailed in the last chapter. Let us now introduce a set of propositional constraints that capture these axioms, called Δ .

Δ contains the following formulae:

- $p_{t=t}$ for every $t \in GT(\Sigma)$
- $p_{t=t'} \Rightarrow p_{t'=t}$, for every $t, t' \in GT(\Sigma)$
- $(p_{t_1=t_2} \wedge p_{t_2=t_3}) \Rightarrow p_{t_1=t_3}$, for every $t_1, t_2, t_3 \in GT(\Sigma)$
- $$\left(\bigwedge_{i \in [1, n]} p_{t_i=t'_i} \right) \Rightarrow \left(p_{R(t_1, \dots, t_n)} \Leftrightarrow p_{R(t'_1, \dots, t'_n)} \right)$$
 for every relation symbol R of arity n .
- $$\left(\bigwedge_{i \in [1, n]} p_{t_i=t'_i} \right) \Rightarrow p_{f(t_1, \dots, t_n) = f(t'_1, \dots, t'_n)}$$
 for every function symbol f of arity n .

It is now easy to show that Γ is satisfiable iff $\Gamma_p \cup \Delta$ is satisfiable. (Proof: If Γ is satisfied in a model M , define a valuation that assigns the propositions p_a to true iff a is true in the model, and argue that Γ_p and Δ will be satisfied under this valuation. Conversely, if $\Gamma_p \cup \Delta$ is satisfied by a propositional valuation, it is easy to see that the equality relation defined by the propositional formulas is a functional congruence over ground terms, and hence defines a Herbrand model of equivalence classes of terms. Interpreting each relation according to the propositional valuation will satisfy the formulas in Γ .

Now, using compactness theorem for propositional logic, we know that $\Gamma_p \cup \Delta$ is satisfiable iff every finite subset of $\Gamma \cup \Delta$ is satisfiable.

Now let us show the required property. If Γ is unsatisfiable, then $\Gamma_p \cup \Delta$ is unsatisfiable, and hence there is a finite subset F of $\Gamma \cup \Delta$ that is satisfiable. Consider $F' = \Gamma_p \cap F$, which is finite. Then the set of FO formulas corresponding to F' in Γ , i.e., the set of formulas whose propositional abstractions are in F' , is unsatisfiable (since $F' \cup \Delta$ is unsatisfiable). Hence there is a finite subset of Γ that is unsatisfiable. \square

The Algorithm

We now continue and finish our recursively enumerable procedure. Recall that we had left off in showing X^* is unsatisfiable, where X^* was obtained by replacing each universally quantified sentence with all possible instantiations of ground terms.

Using the above lemma, we know that X^* is unsatisfiable iff there is some finite subset of X^* that is unsatisfiable.

We can now build a procedure to find such a finite subset. Recall that for any finite subset of quantifier-free formula, there is a decision procedure (that always

halts) whether the set is satisfiable, from the previous chapter. Let's call this decision procedure *DP-QFE* (decision procedure for quantifier-free equality).

8.5.1 The case of finite sets of formulas

We first consider the case when the set of axioms is finite, and hence X is finite. Note that in this case, we can assume the signature is finite too (as the functions/relations not mentioned in the set of formulas clearly do not matter). Note that X is finite, but X^* can be, however, infinite.

Let us consider the following increasing sets that cover X^* . For any $d \in \mathbb{N}$, let T_d denote the ground terms of depth at most d . Formally, these sets are defined recursively as:

- $T_0 = \{c \mid c \text{ is a constant symbol in } \Sigma\}$
- $T_{d+1} = T_d \cup \{f(t_1, \dots, t_n) \mid t_1, \dots, t_n \in T_d, f \text{ is a function symbol of arity } n\}$

Note that $T_i \subseteq T_j$ for any $i \leq j$, and $\bigcup_{i \in \mathbb{N}} T_i$ is the set of all ground terms.

Our procedure is as follows: Given X , a finite set of universal sentences, we do the following:

1. Set $i := 0$;
2. Repeat forever: {
 3. $R := \{\psi[\bar{t} / \bar{x}] \mid \bar{t} \text{ is a tuple of elements in } T_i\}$.
 4. Check if R is satisfiable, by calling *DP-QFE*(R).
If it is not satisfiable, then report X is unsatisfiable and exit (concluding $A \models \varphi$).
 5. Increment i ;
 6. }

The correctness of the algorithm is straightforward to see. If $A \models \varphi$, then $A \cup \{\neg\varphi\}$, and hence X would be unsatisfiable. Hence X^* is unsatisfiable. Hence there is a finite subset of $F \subseteq X^*$ that is unsatisfiable. Let FT be the set of terms mentioned in F ; then FT is finite. Let i be the maximum depth of the terms in FT . Then in iteration i , the algorithm will instantiate X with all terms of depth i , and hence the set R it constructs will be a superset of F , and hence will be unsatisfiable. Hence the decision procedure call to *DP-QFE* will report unsatisfiable, and the algorithm will halt and report $A \models \varphi$.

On the other hand, if $A \not\models \varphi$, then $A \cup \{\neg\varphi\}$, and hence X would be satisfiable. Hence X^* is satisfiable. In each iteration, the algorithm constructs R which is a subset of X^* , and hence the call to *DP-QFE* will report satisfiable in each round. Hence the algorithm will not halt, and will never declare $A \models \varphi$ holds.

8.5.2 The case for infinite sets of formulas

Let us assume the signature is finite. Let us assume we are asked whether $A \models \varphi$, where A is infinite, but recursive. Again, we know that $A \models \varphi$ iff $A \cup \{\neg\varphi\}$ is unsatisfiable iff the set X constructed by converting formulas in the set to universal formulas is unsatisfiable. This set X is unsatisfiable iff X^* is unsatisfiable. And X^* is unsatisfiable iff there is a finite subset of X^* that is unsatisfiable. The key difficulty is to explore larger and larger finite subsets of X^* systematically such that for every finite subset F of X^* , we eventually will explore a superset of F . There are two infinities to consider here—the set of formulas in X is infinite and the set of terms to instantiate the formulas is also infinite. We need to *dovetail* through the two infinities in order to build our procedure.

Let $En : Y_0, Y_1, Y_2, \dots$ be an enumeration of certain finite subsets of X^* . Such an enumeration is said to be *fair* if for every finite subset $F \subseteq_{fin} X^*$, there is some $i \in \mathbb{N}$ such that $F \subseteq Y_i$.

There are several ways to achieve a fair enumeration. We give just one example. Consider the enumeration where Y_i consists of the first i sentences in X enumerated by all possible ground terms of depth i . Then clearly this is a fair enumeration. Let F be a finite subset of X^* . Let i be the largest number such that the i 'th formula in X , instantiated in some way, belongs to F . Let j be the depth of the largest term that was used to instantiate some element in F . Now, let $k = \max(i, j)$. Then it follows that $F \subseteq Y_i$.

For any fair enumeration, we have the following semi-algorithm: Given X , a recursive but infinite set of universal sentences, we do the following:

1. Fix a fair enumeration Y_0, Y_1, \dots of X^*
2. Set $i := 0$;
3. Repeat the following forever: {
4. Check if Y_i is satisfiable, by calling *DP-QFE*(Y_i).
If it is not satisfiable, then report X is unsatisfiable and exit (concluding $A \models \varphi$).
5. Increment i ;
6. }

Again, the proof that the above algorithm always halts when $A \models \varphi$ and reports that it is so, and the proof that when $A \not\models \varphi$, the algorithm runs forever, is easy to see.

We can extend the above argument also to *countably infinite signatures* and infinite but recursive set of axioms. In this case, we need to dovetail between several infinities—exploring more symbols in the signature, exploring more axioms involving this expanding signature, and systematic term instantiation involving this expanding signature. Again, any fair enumeration will give an r.e. procedure.

8.5.3 Completeness Theorem

We can now phrase our completeness theorem, which follows from the above results.

Theorem 8.11 (Completeness) *Let Σ be a finite or countable signature. Let A be a finite set of sentences or an infinite recursive set of sentences over Σ , and let φ be a sentence over Σ . Then the problem of checking whether $A \models \varphi$, is recursively enumerable.*

Let us now work out an example.

Example 8.12 Consider the group axioms, where we have a special constant e for the identity element:

- Associativity: $\forall x, y, z. f(f(x, y), z) = f(x, f(y, z))$
- Identity: $\forall x. f(x, e) = x \wedge f(e, x) = x$
- Inverse: $\forall x \exists y. f(x, y) = e \wedge f(y, x) = e$

Let us now take the above three sentences as the set of axioms A . And let us try to prove the following formula, which says the identity is unique, i.e.,

$$\varphi : \forall e'. ((\forall x. (f(x, e') = x \wedge f(e', x) = x) \Rightarrow (e = e')))$$

Of course, the above property is true even of monoids, i.e., even when the first two axioms hold. However, let's consider all group axioms for this example.

The first two formulas are already in prenex rectified normal form and universal. Skolemizing the third axiom using a new function g gives:

$$\forall x f(x, g(x)) = e \wedge f(g(x), x) = e$$

The new function symbol g intuitively corresponds to a function that provides the inverse of an element. (We don't need to know it is unique in order to ask that such a function exists.)

The formula φ is not in prenex form; bringing it to prenex form gives:

$$\begin{aligned} \varphi &\equiv \forall e'. (\neg(\forall x. (f(x, e') = x \wedge f(e', x) = x) \vee (e = e'))) \\ &\equiv \forall e'. ((\exists x. (\neg(f(x, e') = x) \vee \neg(f(e', x) = x)) \vee (e = e'))) \\ &\equiv \forall e'. \exists x. (\neg(f(x, e') = x) \vee \neg(f(e', x) = x) \vee (e = e')) \end{aligned}$$

The negation of φ is hence:

$$\neg\varphi \equiv \exists e'. \forall x. (f(x, e') = x \wedge f(e', x) = x \wedge \neg(e = e'))$$

Skolemizing the above, by replacing the quantified variable e' by a new constant symbol c gives:

$$\forall x. (f(x, c) = x \wedge f(c, x) = x \wedge \neg(e = c))$$

We now have a set X containing four universal formulas:

- $\forall x, y, z. f(f(x, y), z) = f(x, f(y, z))$
- $\forall x. f(x, e) = x \wedge f(e, x) = x$
- $\forall x. f(x, g(x)) = e \wedge f(g(x), x) = e$
- $\forall x. (f(x, c) = x \wedge f(c, x) = x \wedge \neg(e = c))$

And our task is to check whether they are simultaneously satisfiable.

Let us instantiate with the depth 0 ground terms, i.e., by constants e and c . Then we get the formulae where all quantified variables are replaced by all possible combinations of e and c . That's 14 quantifier-free formulae!

Note that this set includes the following formulae:

- The second formula with x replaced by c :

$$f(c, e) = c \wedge f(e, c) = c$$
- The fourth formula with x replaced by e :

$$(f(e, c) = e \wedge f(c, e) = e \wedge \neg(e = c))$$

Clearly these two formulae are not satisfiable in any model. If $f(c, e) = c$ and $f(c, e) = e$, then we must have $c = e$, which contradicts the conjunction $\neg(e = c)$.

Hence when we ask the decision procedure for quantifier-free formulae whether the 14 formulae have a model, it will report unsatisfiable, and the algorithm above would conclude $A \models \varphi$.

We invite the reader to in fact generate the above formulae, and give them to an SMT solver, like Z3 or CVC4, in order to check that the quantifier-free formulae are unsatisfiable.

Example 8.13 We can take the same axioms above, and try to show that the following holds, which says that inverses are unique. Since we have used the function g , during Skolemization of the axioms, to give us the inverse of elements, let's use the same function g (for brevity).

$$\varphi : \forall x, y. (f(x, y) = e \wedge f(y, x) = e) \Rightarrow (y = g(x))$$

Negating the above and Skolemizing using two new constant symbols c and d gives:

$$(f(c, d) = e \wedge f(d, c) = e) \wedge \neg(d = g(c))$$

Instantiating the Skolemized axioms and the above formula with depth 0 terms (i.e., by the constants e, c , and d) gives a large set of quantifier-free formulae, and it turns out that they are already unsatisfiable. We encourage the reader to write these formulae and feed it to an SMT solver to check that this is indeed so. Consequently, $A \models \varphi$.

8.6 Observations and Consequences

Using SMT solvers:

The above presentation was carefully done so that we get an r.e. procedure that repeatedly calls a solver to check satisfiability of quantifier-free formulae with equality. One can instead also go all the way down to propositional logic satisfiability, and implement the satisfiability of quantifier-free formulae using satisfiability of a propositional encoding of it. This was in fact proposed by Gilmore in 1960! Since SMT solvers already implement satisfiability of quantifier-free formulae with equality, and avoids the blow-up that the propositional encoding entails, we prefer this technique. Furthermore, we will see another application of this term instantiation in a later chapter that allows us to *combine* quantified theories.

The Bernays-Schönfinkel-Ramsey/EPR class

Let us now consider a signature without any function symbols, and a finite set S of formulas of the form $\exists \bar{x} \forall \bar{y} \varphi$. We are asked to check if S is satisfiable. Skolemizing these formulas could introduce new constants but *no new functions*. Consequently, we end up with a set of universal formulas X that we need to check for satisfiability. Since there are no function symbols, the only ground terms are the constants, and we can assume that the constants are only those that occur in the formula, without loss of generality. Consequently, the r.e. procedure outlined earlier in this section can stop after the first instantiation of constants! Hence it is a decision procedure (which always halts on all inputs) and decides satisfiability of such formulae. This fragment of FO formulae, namely $\exists^* \forall^*$ sentences over a signature that has no function symbols, hence admits a decidable satisfiability problem, and is called the *Bernays–Schönfinkel–Ramsey* class or the *effectively propositional reasoning* (EPR) class. Note that for validity, the fragment that is decidable is the $\forall^* \exists^*$ fragment where the signature has no function symbols. This is one of the few quantified fragments of first-order logic that admits decidable validity.

Decidability when Axioms are Negation Complete

A set of axioms A is said to be *consistent* (without contradiction) if there is no sentence such that $A \models \varphi$ and $A \models \neg\varphi$, i.e., $\varphi, \neg\varphi \in Th(A)$. Note that a set of axioms A is consistent iff there is at least one model satisfying the axioms A .

A set of axioms A is said to be *complete* (or *negation complete*) if for every sentence φ , either $A \models \varphi$ or $A \models \neg\varphi$. In other words, the theory of A , $Th(A)$, contains either φ or $\neg\varphi$.

For example, the set of axioms of Presburger arithmetic is consistent and complete. The set of axioms of groups is consistent but not complete.

A consequence of the results of this section is that the theory any complete and consistent axiomatizations is *decidable*. Given a sentence φ , we can execute two copies of the r.e. procedure defined in this section to check whether $A \models \varphi$ and whether $A \models \neg\varphi$. These two executions must be simulated essentially in parallel—for example, running one procedure k steps and then switching to the other for k steps, and then switching back, forever, for some fixed k . Since either $A \models \varphi$ or $A \models \neg\varphi$, one of these procedures will terminate, in which we can halt, and report whether φ is in the theory or not.

Theorem 8.14 *Let A be a recursive set of sentences that is complete. Then the theory of A , $\text{Th}(A)$ is decidable.*

The above also means that if the theory of a single structure is *undecidable*, then it is not axiomatizable. We will prove (see next chapter) that the theory of $(\mathbb{N}, 0, 1, +, \times)$ is undecidable. This means that there is *no recursive set of FO axioms A* such that the theory of A is identical to the theory of this model! This is in fact a version of Gödel's first incompleteness theorem.

Axiomatizability and recursive enumerability

We proved completeness for any recursive set of axioms. However, it is easy to extend the result even when the axioms are recursively enumerable—the procedure will enumerate axioms and instantiate them systematically.

Consider a class of structures C . The notions of having a recursively enumerable set of axioms A that characterize the theory (i.e., $\text{Th}(A) = \text{Th}(C)$ and having $\text{Th}(C)$ itself being recursively enumerable are synonymous. If a r.e. set of axioms A exists characterizing C , then by the completeness theorem, $\text{Th}(\mathcal{A})$ is r.e. as well. On the other hand, if $\text{Th}(C)$ is r.e., then we can choose as axioms this theory itself.

Axiomatic Systems

The most important consequence of the completeness theorem is that it justifies the axiomatic approach. We are typically interested in logic over a particular single structure, or interested in a class of structures. There are many ways to define such a single structure or a class of structures, even using finite means (for example, we can define them using computable functions—giving functions that decide which strings over an alphabet are the elements of a univers, and providing programs that operationally define functions and relations). The axiomatic method, in contrast, asks the class of structures to be defined using properties which are themselves written in FOL. And the completeness theorem gives the guarantee that validity of such a set of axioms is always r.e., which roughly means that every theorem has a proof.

Compactness Theorem for FOL

Another consequence of the results of this section is that the compactness theorem holds for first-order logic sentences as well.

Theorem 8.15 *Let Γ be a set of first-order sentences over a countable signature Σ . Γ is satisfiable iff every finite subset of Γ is satisfiable.*

Proof The forward direction is trivial. For the converse, assume Γ is unsatisfiable. Then, by the results of this section, we can assume Γ is a set of universal sentences. Then, by Lemma 6.3 (Term Expansion Lemma), $\Gamma^* = \{\psi[\bar{t}/\bar{x} \mid \forall x\psi \in \Gamma, t \in GT(\Sigma)]\}$ is unsatisfiable. In other words, the set of quantifier-free sentences obtained by instantiating variables by all possible ground terms is unsatisfiable. By Lemma 6.4, there exists a finite subset F^* of Γ^* that is unsatisfiable. Let $F \subseteq \Gamma$ be a finite subset of Γ from which the elements of F^* were obtained (using term instantiation). Then F is unsatisfiable as well (since even instantiations of variables by ground terms make it unsatisfiable). Hence there is a finite subset of Γ that is unsatisfiable. \square

Appendix A

Computability and Complexity Theory

A Brief Primer

We will now review the basic definitions and theorems in the area of *computational complexity*, which tries to study various models of computation with the goal of understanding their relative computational power, and classify computational problems in terms of computational resources they need. Here, we will primarily consider time and space as the principal resources we will measure for an algorithm.

Recall that the computational problems one studies in the context of theoretical computer science are usually *decision problems*. Decision problems are those where given an input, one expects a Boolean answer. Typically, input instances are encoded as strings over some alphabet of symbols. A decision problem partitions inputs into those for which the expected answer is “yes”/“true” and those for which the answer is “no”/“false”. Therefore, a decision problem is often identified with a *language*, or a collection of strings, namely, those for which the problem demands a “yes” answer. Similarly, the machines we will define, will answer “yes”/“accept” or “no”/“reject” on input strings, and we associate a language $\mathbf{L}(M)$ with machine M , which is the collection of all strings it accepts. Given this interpretation of problems and machines, we will typically say that a machine M solves a problem L (or rather *accepts/recognizes*) if $L = \mathbf{L}(M)$, i.e., M answers “yes” on exactly the inputs that the problem demands the answer to be “yes”.

The main model of computation that we will consider is that of a Turing machine. However before introducing this model, let us recall some of the notation on strings and languages that we will use.

Alphabet, Strings, and Languages.

An *alphabet* Σ is a finite set of elements. A (finite) *string* over Σ is a (finite) sequence $w = a_0a_1 \cdots a_k$ over Σ (i.e., $a_i \in \Sigma$, for all i). The *length* of a string $w = a_0a_1 \cdots a_k$, denoted $|w|$, is the number of elements in it, which in this case is $k + 1$. The unique string of length 0, called the *empty string*, will be denoted by ε . For a string $w = a_0a_1 \cdots a_k$, the i th symbol of the string a_i will be denoted

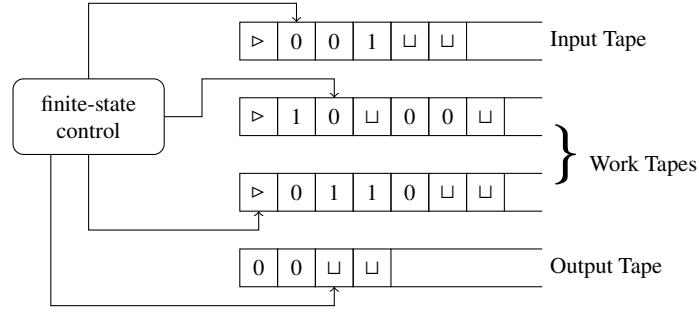


Fig. A.1 Turing machine with a read-only input tape, finitely many read/write worktapes, and a write-only output tape.

as $w[i]$ ¹. For strings $u = a_0a_1 \cdots a_k$ and $v = b_0b_1 \cdots b_m$, their *concatenation* is the string $uv = a_0a_1 \cdots a_kb_0b_1 \cdots b_m$. The set of all (finite) strings over Σ is denoted by Σ^* ; we will sometimes use Σ^i to denote the set of strings of length i . A *language* A is a set of strings, i.e., $A \subseteq \Sigma^*$. Given languages A, B , their concatenation $AB = \{uv \mid u \in A, v \in B\}$. For a language A , $A^0 = \{\epsilon\}$, and A^i denotes the i -fold concatenation of A with itself, i.e., $A^i = \{u_1u_2 \cdots u_i \mid \forall j. u_j \in A\}$. Finally, the *Kleene closure* of a language A , is $A^* = \bigcup_{i \geq 0} A^i$.

A.1 Turing Machines

We now recall the definition of a Turing machine. Since we will use this model to define the time and space bounds during a computation, as well as define computable functions, the most convenient model to consider is that of a multi-tape Turing machine shown in Fig. A.1. Such a model has a read-only *input* tape, a write-only *output* tape and finitely many read/write *work* tape, and a write-only *output* tape. Intuitively, the machine works as follows. Initially, the input string is written out on the input tape, and all the remaining tapes are *blank*. The tape heads are scanning the leftmost cell of each tape, which we will refer to as cell 0. This cell contains a special symbol \triangleright in every tape except the output tape. This is the *left end marker*, which helps the machine realize which cell is the leftmost cell. We will assume these cells are never overwritten by any other symbol, and whenever \triangleright is read on a particular tape, the tape head of the Turing machine will move right. At any given step of the Turing machine does the following. Based on the current state of its finite control, and symbols scanned by each tape head, the machine will change the state of its finite control, write new symbols on each of the work tapes, and move its heads on the input and work tapes either one cell to the left or one cell to the right. During the step, the machine may also choose to write some symbol on its output tape. If

¹ Here we are assuming the the 0th symbol is the “first”.

it writes something on the output tape, then the output tape head moves one cell to the right. If it does not write anything, then the output tape head does not move. We will assume that the machine has two special *halting* states — q_{acc} and q_{rej} — with the property that the machine cannot take any further steps from these states. These are captured in the formal definition of deterministic Turing machines below.

Definition A.1 A *deterministic Turing machine with k -work tapes* is a tuple $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \triangleright)$ where

- Q is a finite set of control states
- Σ is a finite set of input symbols
- $\Gamma \supseteq \Sigma$ is a finite set of tape symbols. We assume that $\{\sqcup, \triangleright\} \subseteq \Gamma \setminus \Sigma$.
- $q_0 \in Q$ is the initial state
- $q_{\text{acc}} \in Q$ is the accept state
- $q_{\text{rej}} \in Q$ is the reject state, with $q_{\text{rej}} \neq q_{\text{acc}}$, and
- $\delta : (Q \setminus \{q_{\text{acc}}, q_{\text{rej}}\}) \times \Gamma^{k+1} \rightarrow Q \times \{-1, +1\} \times (\Gamma \times \{-1, +1\})^k \times (\Gamma \cup \{\varepsilon\})$ is the transition function; here -1 indicates moving the head one position to the left and $+1$ indicates moving the head one position to the right. For $\delta(p, \gamma_0, \gamma_1, \dots, \gamma_k) = (q, d_0, \gamma'_1, d_1, \gamma'_2, d_2, \dots, \gamma'_k, d_k, o)$, for any $i \in \{0, 1, \dots, k\}$, if $\gamma_i = \triangleright$ then $\gamma'_i = \triangleright$ and $d_i = +1$.

We will now formally describe how the Turing machine computes. For this we begin by first identifying information about the Turing machine that is necessary to determine its future evolution. This is captured by the notion of a *configuration*. A single step of a Turing machine depends on all the factors that determine which transition is taken. This clearly includes the control state, and the symbols being read on the input tape and the work tape. However this is not enough. The contents of the work tape change, and what is stored influences what will be read in a future step. Thus we need to know what is stored in each cell of the work tape. Since the input tape is read-only, its contents remain static and so we don't need to carry around its contents. We also need to know the position of each tape head because that determines what is read in this step, how the contents of a tape will change based on the current step, and what will be read in the future as the heads move. Because of all of these observations, a configuration of a Turing machine is taken to be the control state, the position of the input head, the contents of the work tape, and the position of the work tape head. The work tape contents and head position is often represented as a single string where a special marker indicates the head position. These are captured formally by the definition below.

Definition A.2 (Configurations)

A *configuration* c of a Turing machine $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \triangleright)$ is a member of the set $Q \times \mathbb{N} \times (\Gamma^* \{*\} \Gamma \Gamma^* \sqcup^\omega)^k$ ², where we assume that $* \notin \Gamma$ indicates the position of the head. For example, a configuration $c = (q, i, u_1 * a_1 v_1 \sqcup^\omega, u_2 *$

² \sqcup^ω is an *infinite* sequence of blank symbols. Recall that almost all cells contain \sqcup , and so the tape contents are a string of the form $u \sqcup^\omega$, where u is initial portion of the tape containing some non-blank symbols.

$a_2v_2\sqcup^\omega)$ is the configuration of a 2-work tape Turing machine, whose control state is currently q , the input head is scanning cell i , work tape i ($i \in \{1, 2\}$) contains u_i to left of the head, head is scanning symbol a_i and $v_i\sqcup^\omega$ are the contents of cells to the right of the head.

The *initial configuration* (the configuration of the Turing machine when it starts) is $(q_0, 0, * \triangleright \sqcup^\omega, \dots, * \triangleright \sqcup^\omega)$. An *accepting configuration* is a member of the set $\{q_{\text{acc}}\} \times \mathbb{N} \times (\Gamma^* \{*\} \Gamma \Gamma^* \sqcup^\omega)^k$. In other words, it is a configuration whose control state is q_{acc} . A *halting configuration* is a configuration whose control state is either q_{acc} or q_{rej} , i.e., it is a member of the set $\{q_{\text{acc}}, q_{\text{rej}}\} \times \mathbb{N} \times (\Gamma^* \{*\} \Gamma \Gamma^* \sqcup^\omega)^k$.

Having defined configurations, we can formally define how configurations change in a single step of the Turing machine. We begin by defining a function that updates the work tape. For a work tape $u * av\sqcup^\omega$, $\text{upd}(u * av\sqcup^\omega, b, d)$ is the resulting work tape when b is written and the head is moved in direction d . This can be formally defined as

$$\text{upd}(u * av\sqcup^\omega, b, d) = \begin{cases} ub * \sqcup\sqcup^\omega & \text{if } d = +1 \text{ and } v = \varepsilon \\ ub * cv'\sqcup^\omega & \text{if } d = +1 \text{ and } v = cv' \\ u' * cbv\sqcup^\omega & \text{if } d = -1 \text{ and } u = u'c \end{cases}$$

Recall also that for a finite string $w \in \Gamma^*$, $w[i]$ denotes the i th symbol in the string. We can extend this notion to tape contents that are sequences of the form $w\sqcup^\omega$ as follows.

$$w\sqcup^\omega[i] = \begin{cases} w[i] & \text{if } i < |w| \\ \sqcup & \text{otherwise} \end{cases}$$

Definition A.3 (Computation Step)

Consider configurations $c_1 = (q_1, i_1, u_1 * a_1v_1, \dots, u_k * a_kv_k)$ and $c_2 = (q_2, i_2, t_1, \dots, t_k)$ of Turing machine $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \triangleright)$. Let the input string be w . We say $c_1 \xrightarrow{o} c_2$ (machine M moves from configuration c_1 to c_2 in one step and writes o on the output tape) if the following conditions hold. Let $\delta(q_1, w\sqcup^\omega[i_1], a_1, \dots, a_k) = (p, d_0, b_1, d_1, \dots, b_k, d_k)$. Then,

- $q_2 = p$, and $i_2 = i_1 + d_1$,
- for each i , $t_i = \text{upd}(u_i * a_iv_i, b_i, d_i)$

When the output symbol written during a step is not important, we will write $c_1 \mapsto c_2$ to indicate a step from c_1 to c_2 .

Having defined how the configuration of a Turing machine changes in each step, we can define the result of a computation on an input.

Definition A.4 (Computation)

A *computation* of Turing machine M on input w , is a sequence of configurations c_1, c_2, \dots, c_m such that c_1 is the initial configuration of M , and for each i , $c_i \mapsto c_{i+1}$.

Definition A.5 (Acceptance)

An input w is *accepted* by Turing machine M if there is a computation c_1, c_2, \dots, c_m such that c_m is an accepting configuration.

The *language recognized/accepted* by M is $\mathbf{L}(M) = \{w \mid w \text{ is accepted by } M\}$. We say that a language $A \subseteq \Sigma^*$ is *accepted/recognized* by M if $\mathbf{L}(M) = A$.

Definition A.6 (Halting)

A Turing machine M is said *halt* on input w if there is a computation c_1, c_2, \dots, c_m such that c_m is a halting configuration.

The Turing machine model we introduced with an output tape can be used to compute (partial) functions as follows.

Definition A.7 (Function Computation)

The *partial function* computed by a Turing machine M , denoted \mathbf{f}_M , is as follows. If on input w , M has a halting computation $c_1 \xrightarrow{o_1} c_2 \xrightarrow{o_2} \dots \xrightarrow{o_{m-1}} c_m$ then $\mathbf{f}_M(w)$ is defined and equal to $o_1 o_2 \dots o_{m-1}$. On inputs w such that M does not halt, $\mathbf{f}_M(w)$ is undefined.

We say that a (partial) function g is *computable* if there is a Turing machine M such that for every w , $g(w)$ is defined if and only if $\mathbf{f}_M(w)$ is defined, and whenever $g(w)$ is defined, $g(w) = \mathbf{f}_M(w)$.

Most of the time we will be considering Turing machines that accept or recognize languages, rather than those that compute functions. In this context, the symbols written on the output tape don't matter, and so we will often ignore the output tape when describing transitions and computations of such machines.

A.2 Church-Turing Thesis

The Turing machine model introduced in the previous section, is a canonical model to capture mechanical computation. The Church-Turing thesis embodies this statement by saying that anything solvable using a mechanical procedure can be solved using a Turing machine. Our belief in the Church-Turing thesis is based on decades of research in alternate models of computation, which all have turned out to be computationally equivalent to Turing machines. Some of these models include the following.

- Non-Turing machine models: Random Access Machines, λ -calculus, type 0 grammars, first-order reasoning, π calculus, ...
- Enhanced Turing machine models: Turing machines with multiple 2-way infinite tapes, nondeterministic Turing machines, probabilistic Turing machines, quantum Turing machines, ...
- Restricted Turing machine models: Single tape Turing machines, Queue machines, 2-stack machines, 2-counter machines, ...

We will choose to highlight two of these results, that will play a role in our future discussions. The first is the observation that a one work tape Turing machine is computationally as powerful as the multi-work tape model introduced in Definition A.1.

Theorem A.8 For any k work tape Turing machine M , there is a Turing machine with a single work tape $\text{single}(M)$ such that $\mathbf{L}(M) = \mathbf{L}(\text{single}(M))$ and $\mathbf{f}_M = \mathbf{f}_{\text{single}(M)}$ ³.

Proof of Theorem A.8 can be found in any standard textbook and its precise details are skipped. The idea behind the proof is as follows. The single work tape machine $\text{single}(M)$ will simulate the steps of the k -work tape machine M on any input. But in order to simulate M , $\text{single}(M)$ needs to keep track of M 's configuration at each step. That means keeping track of M 's state, its work tape contents, and its tape head. This $\text{single}(M)$ accomplishes by storing M 's state in its own state, and the contents of all k work tapes of M (including the head positions) on the single work tape of $\text{single}(M)$. In general, cell i of the single work tape, stores cell $(i \div k) + 1$ of tape $i \bmod k$; here $i \div m$ denotes the quotient when i is divided by m and $i \bmod m$ denotes the remainder. Then to simulate a single step of M , $\text{single}(M)$ will make multiple passes over its single work tape, to first identify the symbols on each tape read by M to determine the transition to take, and then update the contents of the tape according to the transition.

The second result relates to the nondeterministic Turing machines. The Turing machine model introduced in Definition A.1 is *deterministic*, in the sense that at any given time during the computation of the machine, there is at most one possible transition the machine can take. *Nondeterminism*, on the other hand, is the computational paradigm where the computing device, at each step, may have multiple possible transitions to *choose from*. As a consequence, on a given input the machine may have multiple computations, and the machine is said to accept an input, if any one of these computations leads to an accepting configuration. Formally, we can define a nondeterministic Turing machine as follows.

Definition A.9 A *nondeterministic Turing machine* with k work tapes (and one input tape⁴) is a tuple $M = (Q, \Sigma, \Gamma, \delta, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \rhd)$, where $Q, \Sigma, \Gamma, q_0, q_{\text{acc}}, q_{\text{rej}}, \sqcup, \rhd$ are just like that for deterministic Turing machine, and

$$\delta : (Q \setminus \{q_{\text{acc}}, q_{\text{rej}}\}) \times \Gamma^{k+1} \rightarrow 2^{Q \times \{-1, +1\} \times (\Gamma \times \{-1, +1\})^k}$$

is the transition function. The transition function, given current state and symbols read on the input and work tapes, returns a set of possible next states, direction to move the input head, and symbols to be written and direction to move the head in for each work tape.

The definition of configurations, initial configuration, accepting and halting configurations is the same as in Definition A.2. The definitions of computation step (Definition A.3), computation (Definition A.4), and acceptance and language recognized (Definition A.5) are also the same. Hence we skip defining these formally.

³ For partial functions f and g , we write $f = g$ to indicate that f and g have the same domains (i.e., they are defined for exactly the same elements), and further when $f(x)$ is defined, $f(x) = g(x)$.

⁴ We assume there is no output tape for a nondeterministic Turing machine since such machines are used for function computation.

Every deterministic Turing machine is a special kind of nondeterministic machine, namely, one which has the property that at each time step there is at most one transition enable. One of the important results concerning nondeterministic Turing machines is that the converse is also true, i.e., nondeterministic Turing machines are not more powerful than deterministic Turing machines.

Theorem A.10 *For every nondeterministic Turing machine N , there is a deterministic Turing machine $\text{det}(N)$ such that $\mathbf{L}(N) = \mathbf{L}(\text{det}(N))$.*

A detailed proof of Theorem A.10 is skipped. It can be found in any standard textbook in theory of computation. The broad idea behind the result is the observation that once the length of computation, and the nondeterministic choices at each step are fixed, a deterministic machine can simulate N for that length, on those choices. Thus, the deterministic Turing machine $\text{det}(N)$ simulates N for increasingly longer computations, and for each length, $\text{det}(N)$ will cycle through all possible nondeterministic choices at each step. If any of these computations is accepting for N , then $\text{det}(N)$ will halt and accept.

A.3 Recursive and Recursively Enumerable Languages

The Church-Turing thesis establishes the canonicity of the Turing machine as a model of mechanical computation. The collection of problems solvable on Turing machines is, therefore, worthy of study. Recall that when a Turing machine M is run on an input string w there are 3 possible outcomes — M may (halt and) accept w , M may (halt and) reject w , or M may not halt on w (and therefore not accept). Depending on how a Turing machine behaves we can define two different classes of problems solvable on a Turing machine.

Definition A.11 A language A is *recursively enumerable/semi-decidable* if there is a Turing machine M such that $A = \mathbf{L}(M)$.

A language A is *recursive/decidable* if there is a Turing machine M that halts on *all* inputs and $A = \mathbf{L}(M)$.

Observe that when a problem A is recursive/decidable, it has a special algorithm that solves it and in addition always halts, i.e., on inputs not in A , this algorithm explicitly rejects. Thus, by definition, every recursive language is also recursively enumerable.

Proposition A.12 *If A recursive then A is recursively enumerable.*

We will denote the collection of recursive languages as REC and the collection of all recursively enumerable languages as RE ; thus, Proposition A.12 can be seen as saying that $\text{REC} \subseteq \text{RE}$. The collection of recursive and recursively enumerable languages enjoy some closure properties that are worth recalling.

Theorem A.13 *REC is closed under all Boolean operations while RE is closed under monotone Boolean operations. That is,*

- If $A, B \in \text{RE}$, then $A \cup B$ and $A \cap B$ are also in RE .
- If $A, B \in \text{REC}$, then \overline{A} , $A \cup B$, and $A \cap B$ are all in REC .

Proof We will focus on the two most interesting observations in Theorem A.13; the rest we leave as an exercise for the reader. The first observation we will prove is the closure of RE under union. Let us assume M_A and M_B are Turing machines recognizing A and B , respectively. The computational problem $A \cup B$ asks one to determine if a given input string w belongs to either A or B . We could determine membership in A and B by running M_A and M_B , respectively, but we need to be careful about *how* we run M_A and M_B . Suppose we choose to first run M_A on w and *then* run M_B on w , then we could run into problems. For example, consider the situation where M_A does not halt on w , but $w \in B$. Then, running M_A followed by M_B will never run M_B and therefore never accept, even though $w \in A \cup B$. Switching the order of running M_A and M_B also does not help. What one needs to instead do is, to run M_A and M_B *simultaneously* on w . How does one M_A and M_B at the same time? There are many ways to achieve this. One way is to initially run one step of M_A and then one step of M_B on w from the initial configuration. If either them accept, the algorithm for $A \cup B$ accepts. If not, it will run M_A for two steps, and M_B for two steps, again starting from the respective initial configurations. Again, the algorithm for $A \cup B$ accepts if either simulation accepts. If not the computations of M_A and M_B are increased by one more step, and this process continues, until at some point one of them accepts.

The second result we would like to focus on is the observation that REC is closed under complementation. Let $A \in \text{REC}$ and let M be a Turing machine that halts on all inputs and $\mathbf{L}(M) = A$. The algorithm \overline{M} for \overline{A} , runs M on input w , and if M accepts it rejects and if M rejects then it accepts. Notice that $\mathbf{L}(\overline{M}) = \overline{A}$ only because M halts on all inputs — if M does not halt on (say) w , then $w \in \overline{A}$ but \overline{M} would never accept w ! \square

The following theorem is a useful way to prove that a problem is decidable.

Theorem A.14 A is recursive if and only if A and \overline{A} are recursively enumerable.

Proof If $A \in \text{REC}$ then $\overline{A} \in \text{REC}$ by Theorem A.13. Then both A and \overline{A} are recursively enumerable by Proposition A.12.

Conversely, suppose A and \overline{A} are recognized by M_A and $M_{\overline{A}}$ respectively. The recursively algorithm M for A , on a given input w , will run both M_A and $M_{\overline{A}}$ simultaneously (as in the proof of Theorem A.13), and accept if either M_A accepts or $M_{\overline{A}}$ rejects. Notice, that any given input w belongs to either A or \overline{A} , and therefore at least one out of M_A and $M_{\overline{A}}$ is guaranteed to halt on each input. Therefore M will always halt. \square

Encodings.

Every object (graphs, programs, Turing machines, etc.) can be encoded as a binary string. The details of the encoding scheme itself are not important, but it should be

simple enough that the data associated with the object should be easily recoverable by reading the binary encoding. For example, one should be able to reconstruct the vertices and edges of a graph from its encoding, or one should be able to reconstruct the states, transitions, etc. of a Turing machine from its encoding. For a list of objects O_1, O_2, \dots, O_n , we will use $\langle O_1, O_2, \dots, O_n \rangle$ to denote their binary encoding. In particular, for a Turing machine M , $\langle M \rangle$ is its encoding as binary string. Conversely, for a binary string x , M_x denotes the Turing machine whose encoding is the string x .

Once we establish an encoding scheme, we can construct a *Universal Turing machine*, which is an *interpreter* that given an encoding of a Turing machine M and an input w , can simulate the execution of M on the input string w . This is an extremely important observation that establishes the recursive enumerability of the membership problem for Turing machines.

Theorem A.15 *There is a Turing machine U (called the universal Turing machine) that recognizes the language $\text{MP} = \{\langle M, w \rangle \mid w \in \mathbf{L}(M)\}$. In other words, $\text{MP} \in \text{RE}$.*

Not every decision problem/language is recursively enumerable. Using Cantor's diagonalization technique, one can establish the following result.

Theorem A.16 *The language $\bar{K} = \{x \mid x \notin \mathbf{L}(M_x)\}$ is not recursively enumerable.*

Proof The proof of Theorem A.16 relies on a diagonalization argument to show that the language of every Turing machine differs from \bar{K} , and therefore \bar{K} is not recursively enumerable.

Consider an arbitrary Turing machine M_x whose encoding as a binary string is x . We will show that $\mathbf{L}(M_x) \neq \bar{K}$, thereby proving the theorem. Observe that if $x \in \mathbf{L}(M_x)$ then by definition $x \notin \bar{K}$ and if $x \notin \mathbf{L}(M_x)$ then again by definition $x \in \bar{K}$. Therefore $x \in (\bar{K} \setminus \mathbf{L}(M_x)) \cup (\mathbf{L}(M_x) \setminus \bar{K}) \neq \emptyset$. \square

A.4 Reductions

Theorem A.16 is the first result that establishes that there are problems that are computationally difficult. Further results on the computational hardness of problems are usually established using the notion of *reductions*. Reductions demonstrate how one problem can be converted into another in such a way that a solution to the second problem can be used to solve the first. Formally, it is defined as follows.

Definition A.17 A (many-one/mapping) reduction from A to B is a computable (total) function $f : \Sigma^* \rightarrow \Sigma^*$ such that for any input string w ,

$$w \in A \text{ if and only if } f(w) \in B$$

In this case, we say A is (many-one/mapping) reducible to B and we denote it by $A \leq_m B$.

Since many-one/mapping reductions are the only form of reduction we will study, we will drop the adjective “many-one” and “mapping” and simply call these reductions. Let us look at a couple of examples of reductions.

Example A.18 Let us consider the complement of MP , i.e., $\overline{\text{MP}} = \{\langle M, w \rangle \mid w \notin \mathbf{L}(M)\}$. One can show that $\overline{K} \leq_m \overline{\text{MP}}$ as follows. The reduction f is the following function: $f(x) = \langle M_x, x \rangle$.

To prove that f is a reduction, we need to argue two things. First that f is computable, i.e., we need to come up with a Turing machine M_f that always halts and produces the string $f(x)$ on input x . In this example, to construct $f(x)$, we simply need to “copy” the string x which clearly is a computable function. Second we need to argue that $x \in \overline{K}$ iff $f(x) \in \overline{\text{MP}}$. This can be argued as follows: $x \in \overline{K}$ iff $x \notin \mathbf{L}(M_x)$ (definition of \overline{K}) iff $\langle M_x, x \rangle \in \overline{\text{MP}}$ (definition of $\overline{\text{MP}}$) iff $f(x) \in \overline{\text{MP}}$ (definition of f).

Example A.19 Consider the problem

$$\overline{\text{HP}} = \{\langle M, w \rangle \mid M \text{ does not halt on } w\}.$$

We will prove that $\overline{K} \leq_m \overline{\text{HP}}$.

Given a binary string x , let us consider the following program H_x .

```

 $H_x(w)$ 
  result =  $M_x(x)$ 
  if (result = accept)
    return accept (* on input w *)
  else
    while true do

```

In other words, the program H_x on input w , ignores its input and runs the program M_x on x . If M_x halts and accepts x then H_x halts and accepts w . Otherwise, H_x does not halt. Thus, the program H_x halts on some (all) inputs if and only if $x \in \mathbf{L}(M_x)$.

Let us now describe the reduction from \overline{K} to $\overline{\text{HP}}$: $f(x) = \langle H_x, x \rangle$. Observe first that f satisfies the properties of a reduction because $x \in \overline{K}$ iff $x \notin \mathbf{L}(M_x)$ iff H_x does not halt on x (and all input strings) iff $\langle H_x, x \rangle \in \overline{\text{HP}}$. To establish that f is a reduction, we also need to argue that f is computable. On input string x , we need a program that produces the source code for H_x (given above) and copies the string x after the source code. This is clearly computable.

Reductions are a way for one to compare the computational difficulty of problems — if A reduces to B then A is at most as difficult as B , or B is at least as difficult as A . This is formally captured in the following proposition.

Theorem A.20 *If $A \leq_m B$ and B is recursively enumerable (recursive) then A is recursively enumerable (recursive).*

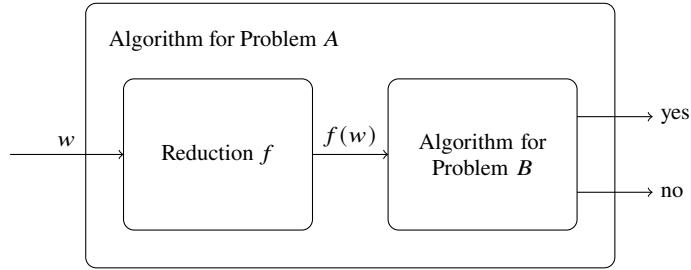


Fig. A.2 Schematic argument for Theorem A.20.

Proof Let f be a reduction from A to B that is computed by Turing machine M_f , and let M_B be a Turing machine that recognizes B . The algorithm for A is schematically shown in Fig. A.2 — on input w , compute $f(w)$ using M_f and run M_B on $f(w)$. Notice that this algorithm always halts if M_B always halts. Thus, if B is recursive then A is also recursive. \square

Theorem A.20 can be seen to informally say “if A reduces to B and B is computationally easy then A is computationally easy”. It is often used in the contrapositive sense and it is useful to explicitly state this observation.

Corollary A.21 *If $A \leq_m B$ and A is not recursively enumerable (undecidable) then B is not recursively enumerable (undecidable).*

We can use the above corollary to argue the computational hardness of some problems.

Theorem A.22 $\overline{\text{MP}}$ is not recursively enumerable. Therefore, MP is undecidable.

Proof Example A.18 establishes that $\overline{K} \leq_m \overline{\text{MP}}$. Together with Theorem A.16 and Corollary A.21, we can conclude that $\overline{\text{MP}}$ is not recursively enumerable. Finally, since $\overline{\text{MP}}$ is not recursively enumerable, Theorem A.14 establishes that MP is not decidable/recursive. \square

Since $\text{MP} \in \text{RE}$ (Theorem A.15) and $\overline{\text{MP}} \notin \text{RE}$ (Theorem A.22), we have a witness to the fact that RE is not closed under complementation. Just like Theorem A.22, we could establish similar properties for the *halting problem*.

Theorem A.23 $\overline{\text{HP}}$ is not recursively enumerable. Therefore, $\text{HP} = \{\langle M, w \rangle \mid M \text{ halts on } w\}$ is undecidable.

Proof Follows from Example A.19 and the argument in the proof of Theorem A.22. \square

Reductions are transitive and hence a *pre-order*; thus, the use of \leq to denote them is justified.

Theorem A.24 The following properties hold for reductions.

- If $A \leq_m B$ then $\overline{A} \leq_m \overline{B}$.
- If $A \leq_m B$ and $B \leq_m C$ then $A \leq_m C$.

Proof If f is a reduction from A to B , then one can argue that f is also a reduction from A to \overline{B} . And, if f is a reduction from A to B and g a reduction from B to C then $g \circ f$ is a reduction from A to C . Establishing these observations to prove the theorem is left as an exercise. \square

Having found a lens to compare the computational difficulty of two problems (namely, reductions), one can use them to argue that a problem is at least as difficult as a whole collection of problems, or something is the “hardest” problem in a collection. This leads us to notions of hardness and completeness.

Definition A.25 A language A is **RE-hard** if for every $B \in \text{RE}$, $B \leq_m A$.

A language A is **RE-complete** if A is **RE-hard** and $A \in \text{RE}$.

Thus, an **RE-complete** problem is the hardest problem that is recursively enumerable, while an **RE-hard** problem is something that is at least as hard as any other **RE** problem. Are there examples of such problems? It turns out that **MP**, **HP**, and **K** are all **RE-complete**. We establish this for **MP** in the following theorem.

Theorem A.26 **MP** is **RE-complete**.

Proof Membership in **RE** has been established in Theorem A.15. So all we need to prove is the hardness. Let B be any recursively enumerable language, and let M be a Turing machine recognizing B . The reduction from B to **MP** is as follows: $f(w) = \langle M, w \rangle$. It is easy to see that $w \in B$ iff $w \in \mathbf{L}(M)$ (since M recognizes B) iff $\langle M, w \rangle \in \mathbf{MP}$ (definition of **MP**) iff $f(w) \in \mathbf{MP}$ (definition of f). It is also easy to see that f is computable — in order to compute $f(w)$, all we need to do is prepend the source code of M . \square

Establishing **RE-hardness** of a problem is sufficient to guarantee it’s undecidability.

Theorem A.27 If A is **RE-hard** then A is undecidable.

Proof If A is **RE-hard** then since **MP** $\in \text{RE}$, we have $\mathbf{MP} \leq_m A$. Since **MP** is undecidable (Theorem A.22), by properties of a reduction (Corollary A.21) A is undecidable. \square

A.5 Complexity Classes

Computational resources needed to solve a problem depend on the size of the input instance. For example, it is clearly easier to compute the sum of two one digit numbers as opposed to adding two 15 digit numbers. The resource requirements of an algorithm/Turing machine are measured as a function of the input size. We will only study time and space as computational resources in this presentation. We

begin by defining time bounded and space bounded Turing machines, which are defined with respect to bounds given by functions $T : \mathbb{N} \rightarrow \mathbb{N}$ and $S : \mathbb{N} \rightarrow \mathbb{N}$ that are non-decreasing, i.e., for all $n \leq m \in \mathbb{N}$, $T(n) \leq T(m)$ and $S(n) \leq S(m)$. Our definitions apply to both deterministic and nondeterministic machines.

Definition A.28 A (deterministic/nondeterministic) Turing machine M is said to run in *time* $T(n)$ if on any input u , *all* computations of M on u take at most $T(|u|)$ steps; here $|u|$ refers to the length of input u .

A (deterministic/nondeterministic) Turing machine M is said to use *space* $S(n)$ if on any input u , *all* computations of M on u use at most $S(|u|)$ work tape cells. In this context, a work tape cell is said to be used if it is written to at least once during the computation. Notice that, if a work tape cell is written multiple times during a computation, it counts as only one cell when measuring the space requirements; thus, work tape cells can be reused without adding to the space bounds.

It is worth examining Definition A.28 carefully. Our requirement for a Turing machine running within some time or space bound applies to *all* computations, whether they are accepting or not. Notice also that the definition is the same for both deterministic and nondeterministic models — in a deterministic machine the *unique* computation on a given input must satisfy the resource bounds, and in a nondeterministic machine, *all* computations on the input must satisfy the bounds. In particular, if a Turing machine (deterministic or nondeterministic) runs within a time bound, then it *halts* in every computation of every input.

Having defined time and space bounded machines, we can define the basic complexity classes which are collections of (decision) problems that can be solved within certain time and space bounds.

Definition A.29 We define the following basic complexity classes.

- A language $A \in \text{DTIME}(T(n))$ iff there is a *deterministic* Turing machine that runs in time $T(n)$ such that $A = \mathbf{L}(M)$.
- A language $A \in \text{NTIME}(T(n))$ iff there is a *nondeterministic* Turing machine that runs in time $T(n)$ such that $A = \mathbf{L}(M)$.
- A language $A \in \text{DSPACE}(S(n))$ iff there is a *deterministic* Turing machine that uses space $S(n)$ such that $A = \mathbf{L}(M)$.
- A language $A \in \text{NSPACE}(S(n))$ iff there is a *nondeterministic* Turing machine that uses space $S(n)$ such that $A = \mathbf{L}(M)$.

Our computational model of Turing machines, and our definitions of time and space bounded computations are robust with respect to constant factors. This observation is captured by two central results in theoretical computer science, namely, the linear speedup and compression theorems. It says that one can always improve the running time or space requirements for solving a problem by a constant factor.

Theorem A.30 (Linear Speedup)

If $A \in \text{DTIME}(T(n))$ (or $A \in \text{NTIME}(T(n))$) and $c > 0$ is any constant, then $A \in \text{DTIME}(cT(n) + n)$ ($A \in \text{NTIME}(cT(n) + n)$).

Proof (Sketch) Let $A = \mathbf{L}(M)$. We will describe a machine M' which will simulate k steps of M in 8 steps; if $k > \frac{8}{c}$, we will get the desired result. M' will have one more work tape, a much larger tape alphabet, and control states than M .

- M' copies the input onto the additional work tape in compressed form: k successive symbols of M will be represented by one symbol in M' . Time taken is n . M' will maintain M 's work tape contents in compressed form on the second work tape as well.
- M' uses the additional work tape as “input tape”. The head positions of M , within the k symbols represented by current cells, is stored in finite control.

One *basic move* of M' (consisting of 8 steps), will simulate k steps of M as follows.

- At the start of basic move, M' moves its tape heads one cell left, two cells right and one cell left, storing the symbols read in the finite control. Now, M' knows all symbols within the radius of k cells of any of M 's tape heads. This takes 4 steps.
- Based on the transition function of M , M' can compute the effect of the next k steps of M .
- Using any additional (at most) 4 steps, M' updates the contents of its tapes as a result of the k steps, and moves the heads appropriately. \square

Theorem A.31 (Linear Compression)

If $A \in \mathbf{DSPACE}(S(n))$ (or $A \in \mathbf{NSPACE}(S(n))$) and $c > 0$ is any constant then $A \in \mathbf{DSPACE}(cS(n))$ ($A \in \mathbf{NSPACE}(cS(n))$).

Proof Increase the tape alphabet size and store work tape contents in compressed form as in Theorem A.30. \square

Theorems A.30 and A.31 suggest that when analyzing the time and space requirements of an algorithm we can ignore constant terms. This leads to the use of the order notation.

Definition A.32 Consider functions $f : \mathbb{N} \rightarrow \mathbb{N}$ and $g : \mathbb{N} \rightarrow \mathbb{N}$.

- $f(n) = O(g(n))$ if there are constants c, n_0 such that for $n > n_0$, $f(n) \leq cg(n)$. $g(n)$ is an *asymptotic upper bound*.
- $f(n) = \Omega(g(n))$ if there are constants c, n_0 such that for $n > n_0$, $f(n) \geq cg(n)$. $g(n)$ is an *asymptotic lower bound*.

The complexity classes identified in Definition A.29 are a very fine classification of problems. They include complexity classes whose classification of problems is sensitive to our use of Turing machines as a model of computation. Ideally we would like to study complexity classes such that if a problem is classified in a certain class then that classification should be “platform independent”. That is, whether we choose to study complexity on Turing machines or Random Access Machines, our observations should still hold. They should also be invariant under small changes to the Turing machine model, like changing the number of work tapes, alphabet, nature

of the tapes, etc. There is a strengthening of the Church-Turing thesis, called the *invariance thesis* articulated by Church, that underlies our belief in the robustness of the Turing machine model, subject to small changes in the time and space bounds. It says that

Any effective, mechanistic procedure can be simulated on a Turing machine using the same space (if space is $\geq \log n$) and only a polynomial slowdown (if time $\geq n$)

In addition to the requirement that complexity classes be robust to changes to the computational platform, we would like the classes to be closed under function composition — making function/procedure calls to solve sub-problems is a standard algorithmic tool, and we would like the complexity to remain the same as long as the sub-problems being solved are equally simple. Finally, we would like our complexity classes to capture natural, “interesting”, real-world problems. For these reasons, we typically study the following complexity classes that provide a coarser classification of problems than that provided in Definition A.29.

Definition A.33 Commonly studied complexity classes are the following.

$$\begin{array}{ll} L = \text{DSPACE}(\log n) & NL = \text{NSPACE}(\log n) \\ P = \cup_k \text{DTIME}(n^k) & NP = \cup_k \text{NTIME}(n^k) \\ \text{PSPACE} = \cup_k \text{DSPACE}(n^k) & \text{NPSPACE} = \cup_k \text{NSPACE}(n^k) \\ \text{EXP} = \cup_k \text{DTIME}(2^{n^k}) & \text{NEXP} = \cup_k \text{NTIME}(2^{n^k}) \end{array}$$

In addition to the above classes, for any class C , $\text{co}C = \{A \mid \overline{A} \in C\}$. Please note that $\text{co}C$ is *not* the complement of C but instead is the collection of problems whose complement is in C .

A.6 Relationship between Complexity Classes

We begin by relating time and space complexity classes.

Theorem A.34 $\text{DTIME}(T(n)) \subseteq \text{DSPACE}(T(n))$ and $\text{NTIME}(T(n)) \subseteq \text{NSPACE}(T(n))$

Proof A Turing machine can scan at most one new work tape cell in any step. Therefore, the number of work tape cells used during a computation cannot be more than the number of steps. \square

Theorem A.35 $\text{DSPACE}(S(n)) \subseteq \text{DTIME}(n \cdot 2^{O(S(n))})$ and $\text{NSPACE}(S(n)) \subseteq \text{NTIME}(n \cdot 2^{O(S(n))})$. In particular, when $S(n) \geq \log n$, we have $\text{DSPACE}(S(n)) \subseteq \text{DTIME}(2^{O(S(n))})$ and $\text{NSPACE}(S(n)) \subseteq \text{NTIME}(2^{O(S(n))})$.

We will skip giving a direct proof of Theorem A.35. It will follow from Theorem A.37 and Theorem A.38. An immediate consequence of Theorems A.34 and A.35 are the following relationships between the complexity classes.

Corollary A.36

$$\begin{aligned} L \subseteq P \subseteq \text{PSPACE} \subseteq \text{EXP} \\ NL \subseteq NP \subseteq \text{NPSPACE} \subseteq \text{NEXP} \end{aligned}$$

We will establish relationships between deterministic and nondeterministic complexity classes.

Theorem A.37 $\text{DTIME}(T(n)) \subseteq \text{NTIME}(T(n))$ and $\text{DSPACE}(S(n)) \subseteq \text{NSPACE}(S(n))$.

Proof This follows from the fact that, by definition, every deterministic Turing machine is a special nondeterministic Turing machine, namely, those that have exactly one transition enabled from every non-halting configuration. \square

Nondeterministic complexity class can also be related to deterministic complexity classes. In fact, we now prove a result that subsumes the containment results for nondeterministic classes established in Theorems A.34 and A.35.

Theorem A.38 $\text{NTIME}(T(n)) \subseteq \text{DSPACE}(T(n))$ and $\text{NSPACE}(S(n)) \subseteq \text{DTIME}(n2^{O(S(n))})$.

Proof Let us begin by proving the first inclusion. Consider $A \in \text{NTIME}(T(n))$ and let M be $T(n)$ -time bounded nondeterministic machine recognizing A . On any input w of length n , the computations of M can be organized as a tree, and since M runs in time $T(n)$, this tree has height $T(n)$. Now the deterministic algorithm D to solve A will perform a *depth first search* (DFS) on this computation tree of M , constructing this tree as it is explored, and accepting if some node in this computation tree corresponds to an accepting configuration. The space needed by D to perform this DFS is the memory needed to store the call stack. The stack during a DFS keeps track of the path being currently explored in the tree to enable backtracking. Since the computation tree of M is of height $T(n)$, the height of the call stack is also $T(n)$. A naïve implementation of the DFS algorithm will store the sequence of tree vertices on the current path; since in this case each vertex is a configuration of M , these can be represented by strings of length $T(n)$ (as work tape cells cannot exceed $T(n)$ as in Theorem A.34). This gives us a space bound of $T(n)^2$ for algorithm D . However, instead of storing the actual configurations in the computation being currently explored, D can just store the sequence of nondeterministic choices made by M in the current computation. With this information about the nondeterministic choices, D can *reconstruct* the configuration at the end of a sequence of steps, by resimulating M from the beginning — this increases the running time of D , but reduces the space requirements of D which is what we care about for this result. If M has k choices at each step, the stack of D during DFS is simply a k -ary string of length $\leq T(n)$, which means that D is $T(n)$ -space bounded.

For the second result, let us consider $A \in \text{NSPACE}(S(n))$ and a nondeterministic Turing machine M that recognizes A in $S(n)$ space. On a given input w of the length n , it is useful to define the notion of a *configuration graph* of M . The configuration graph is a directed graph that has as vertices, configurations of M , and has an edge from c_1 to c_2 , if M can move from configuration c_1 to configuration c_2 in one step

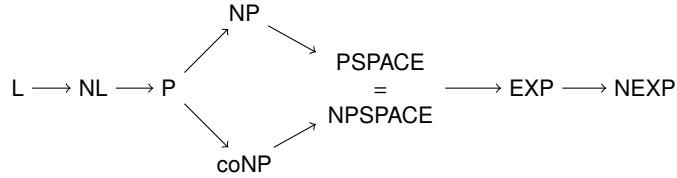


Fig. A.3 Relationship between Complexity Classes. \rightarrow indicates containment, though whether it is strict is unknown.

given input w . Observe that M accepts w if an accepting configuration is reachable from the initial configuration in this configuration graph. Notice also that since M is $S(n)$ -space bounded, the total number of vertices in this graph is $\leq n2^{O(S(n))}$ (see proof of Theorem A.35). Now we can run our favorite graph search algorithm (depth first search or breadth first search) on this configuration graph to see if an accepting configuration is reachable; the graph will be constructed on-the-fly as it is being explored. Such an algorithm (which deterministic), takes time that is linear in the size of the graph, which gives us a $n2^{O(S(n))}$ deterministic algorithm for A . \square

Our new observations relating deterministic and nondeterministic complexity classes gives us the following relationships.

Corollary A.39

$$L \subseteq NL \subseteq P \subseteq NP \subseteq PSPACE \subseteq NPSPACE \subseteq EXP \subseteq NEXP$$

An important result due to Savitch, relates nondeterministic and deterministic space complexity classes. The interested reader can find its proof in textbooks like [?].

Theorem A.40 (Savitch)

For $S(n) \geq \log n$, $NPSPACE(S(n)) \subseteq DSPACE(S(n)^2)$. In particular, this means that $PSPACE = NPSPACE$.

Putting all our observations together we get the relationships shown in Fig. A.3. It is worth observing that for any deterministic complexity class C , $C = \overline{co}C$; this is because an algorithm for \overline{A} is to run the deterministic algorithm for A and flip the final answer. Thus, $P = \overline{co}P$. Further, from Theorem A.37, we have $\overline{co}P \subseteq \overline{co}NP$, giving us the containment $P \subseteq \overline{co}NP$. Finally, due the space hierarchy theorem, we also know that $L \neq PSPACE$ and $NL \neq PSPACE$, and from the time hierarchy theorem, we know that $P \neq EXP$ and $NP \neq NEXP$; the hierarchy theorems are beyond the scope of this brief primer.

A.7 P and NP

Cobham-Edmonds Thesis, named after Alan Cobham and Jack Edmonds, asserts that the only computational problems that have “efficient” or “feasible” algorithmic

solutions are those that belong to P . In other words, P is the collection of *tractable* computational problems. There are many features of P that justify this view.

- The invariance thesis suggests that any problem in P can be solved in polynomial time on *any* reasonable computational model. Thus, the statement of a problem being efficiently computable is platform independent.
- Most encodings of an input structure are polynomially related in terms of their length. Thus, if a problem is in P for one encoding, it will be in P even if input instances are encoded in a different manner. Therefore, P is insensitive to problem encodings.
- Most natural problems in P have algorithms whose running time is bounded by a low-order polynomial. Thus, their running times are likely to be low for most problem instances.
- The asymptotic growth of polynomials is moderate when compared to the astounding growth of exponential functions. Thus, problems in P are likely to be feasibly solved even on large problem instances.

The crux of the Cobham-Edmonds thesis is that for problem to be solvable in practice, it should have a polynomial time algorithm. Therefore, much effort in the past 50 years has been devoted to understanding the class of problems in P . In particular, can we prove that the complexity classes containing P in Fig. A.3, like NP and PSPACE , also have efficient solutions, i.e., are contained in P ? Or can we say for certain that some problems in NP and PSPACE *cannot* be solved efficiently?

A.7.1 Alternate characterization of NP

We defined NP as the collection of problems that can be solved in polynomial time on a nondeterministic Turing machine. In this section, we will give an alternate definition, namely, as those problems that are efficiently “verifiable”.

Definition A.41 A language A is *polynomially verifiable* if there is a $k \in \mathbb{N}$ and a deterministic Turing machine V such that

$$A = \{w \mid \exists p. V \text{ accepts } \langle w, p \rangle\}$$

and V takes at most $|w|^k$ steps on input $\langle w, p \rangle$, i.e., V running time is *independent* of the length of p . Here V is called a *verifier* for A , and for $w \in A$, the strings p such that $\langle w, p \rangle$ is accepted by V are called a *proof* of w (with respect to V).

The notion of a language A being polynomial verifiable says that when a string $w \in A$, there is a *proof* p (maybe even more than one) such that w augmented with p “convinces” V , i.e., causes V to accept. However, if $w \notin A$ then there is no proof string p that can convince V of w ’s membership in A . Notice also that V ’s running time on input $\langle w, p \rangle$ is independent of the length of p ; it always runs in time $|w|^k$ no matter what p is. Now, since in $|w|^k$ steps, V cannot read more than $|w|^k - |w|$ bits of p , we can without loss of generality assume that p is a string whose length

is bounded by a polynomial in the length of w . Thus, we could informally say that A is polynomially verifiable, if for any string $w \in A$ there is a “short” proof (of polynomial length) that can be efficiently checked (in polynomial time) by a verifier, and if $w \notin A$ there is no proof that can convince a verifier.

A language being polynomially verifiable is equivalent to a problem having a nondeterministic polynomial time verifier.

Theorem A.42 $A \in \text{NP}$ if and only if A is polynomially verifiable.

Proof Consider $A \in \text{NP}$, and let M be nondeterministic Turing machine recognizing A in time n^k for some k . We can assume without loss of generality that M has at most two choices at any given step. The verifier V for A will work as follows. On input $\langle w, p \rangle$, where p is a binary string, it will first copy w onto a work-tape, and compute n^k . It will then simulate M for n^k steps using the work-tape with w as the input tape, taking p to be the sequence of nondeterministic choices. V accepts $\langle w, p \rangle$ if M accepts w with p as the nondeterministic choices. Observe that V is a deterministic algorithm running in $O(n^k)$ time on $|w| = n$. Further $A = \{w \mid \exists p. V \text{ accepts } \langle w, p \rangle\}$.

Conversely, suppose V is a polynomial time verifier for A . Suppose V runs in time $|w|^k$ on input $\langle w, p \rangle$. The nondeterministic algorithm M for A will work as follows. On input w , M will guess a string p of length $|w|^k$. Then M will simulate V on w and the guessed string p , accepting if and only if V accepts. It is easy to see that $L(M) = A$ and M runs in time $O(n^k)$. \square

Thus, NP is the collection of all problems A whose membership question has short, efficiently checkable proofs. The question of whether all problems in NP have polynomial time algorithms — whether $\text{P} \stackrel{?}{=} \text{NP}$ — is thus the question of whether every problem that has a short, efficiently checkable proofs also have the property that these proofs can be *found* efficiently. Phrased in this manner, the likely answer seems to be no. There are also results that seem to suggest that P is likely to be not equal to NP , though a firm resolution of this question has eluded researchers for the past 50 years.

A.7.2 Reductions, Hardness and Completeness

In an effort to resolve the P versus NP question, researchers have tried to identify canonical problems whose study can help address this challenge. The goal is to identify, in some sense, the most difficult problems in NP such that either (a) they are candidate problems that may not have polynomial time algorithms, or (b) finding a polynomial time algorithms for these problems will constructively demonstrate that $\text{P} = \text{NP}$. In order to identify such difficult problems, we need to be able to compare the difficulty of two problems. For this the most convenient technique is that of reductions. Unlike, many-one reductions introduced before, we will require that these be computed in polynomial time.

Definition A.43 A *polynomial time reduction* from A to B is a *polynomial time computable function* f such that for every input w ,

$$w \in A \text{ if and only if } f(w) \in B.$$

In such a case we say that A is *polynomial time reducible* to B and is denoted by $A \leq_P B$.

Example A.44 Consider the following problems.

$$\begin{aligned} \text{SAT} &= \{\langle \varphi \rangle \mid \varphi \text{ is in CNF and } \varphi \text{ is satisfiable}\} \\ k\text{-COLOR} &= \{\langle G, k \rangle \mid G \text{ is an undirected graph that can be colored using } k \text{ colors}\} \end{aligned}$$

Proposition 1.35 shows that for any graph G and $k \in \mathbb{N}$ there is a set of clauses $\Gamma_{G,k}$ such that G is k -colorable if and only if $\Gamma_{G,k}$ is satisfiable. The number of clauses in $\Gamma_{G,k}$ is proportional to the number of vertices and edges in G , and each clause has at most k -literals. It is also easy to see that $\Gamma_{G,k}$ can be constructed from G in time that is linear in the size of G . Thus, these observations together establish that $k\text{-COLOR} \leq_P \text{SAT}$.

Example A.45 A formula φ in CNF is said to be in 3-CNF if every clause in φ has exactly 3 literals. For example, $(x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_4) \wedge (x_4) \wedge (\neg x_2 \vee \neg x_3 \vee x_4)$ is not in 3-CNF, while $(x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_4 \vee x_2)$ is in 3-CNF. Recall the **SAT** problem is one whether given a formula φ in CNF, we need to determine if φ is satisfiable. A special case of this problem is one where the input formula is promised to be in 3-CNF. Formally we have,

$$3\text{-SAT} = \{\langle \varphi \rangle \mid \varphi \text{ is in 3-CNF and is satisfiable}\}.$$

Since 3-SAT is a “special” version of SAT, the identity function is a reduction from 3-SAT to SAT; thus, $3\text{-SAT} \leq_P \text{SAT}$. It turns out the one also has a reduction the other way around.

The reduction from SAT to 3-SAT is as follows. Consider a CNF formula φ ; it will be convenient to think of φ as a set of clauses. Our reduction will convert (in polynomial time) each clause $c \in \varphi$ into a 3-CNF formula $f(c)$ such that c and $f(c)$ are satisfied by (almost) the same set of truth assignments. Then $f(\varphi) = \{f(c) \mid c \in \varphi\}$, and it will be the case that φ is satisfiable iff $f(\varphi)$ is satisfiable.

Let us now describe the translation of clauses. The translation of clause c will depend on how many literals c has. Let $c = \ell_1 \vee \ell_2 \vee \cdots \vee \ell_k$. Depending on k , we have the following cases.

Case $k = 1$ Let u and v be “new” propositions not used before. Define $f(c)$ to be

$$(\ell_1 \vee u \vee v) \wedge (\ell_1 \vee u \vee \neg v) \wedge (\ell_1 \vee \neg u \vee v) \wedge (\ell_1 \vee \neg u \vee \neg v)$$

Case $k = 2$ Let u be a “new” proposition. $f(c)$ is given by

$$(\ell_1 \vee \ell_2 \vee u) \wedge (\ell_1 \vee \ell_2 \vee \neg u)$$

Case $k = 3$ In this case $f(c) = c$.

Case $k > 3$ Let y_1, y_2, \dots, y_{k-3} be new propositions. Then $f(c)$ is

$$\begin{aligned} & (\ell_1 \vee \ell_2 \vee y_1) \wedge (\ell_3 \vee \neg y_1 \vee y_2) \wedge (\ell_4 \vee \neg y_2 \vee y_3) \wedge \dots \\ & \wedge (\ell_{k-2} \vee \neg y_{k-4} \vee y_{k-3}) \wedge (\ell_{k-1} \vee \ell_k \vee \neg y_{k-3}) \end{aligned}$$

It is easy to see that f can be computed in time that is linear in the size of φ . Moreover, φ is satisfiable iff $f(\varphi)$ is satisfiable (left as exercise). Thus, f is a polynomial time reduction showing $\text{SAT} \leq_P 3\text{-SAT}$.

Polynomial time reductions satisfy properties similar to many-one reductions: they are transitive and if A reduces to B then \overline{A} reduces to \overline{B} .

Proposition A.46 *The following properties hold for polynomial time reductions.*

- If $A \leq_P B$ then $\overline{A} \leq_P \overline{B}$.
- If $A \leq_P B$ and $B \leq_P C$ then $A \leq_P C$.

Proof Detailed proof of these observations is left as an exercise. But the sketch is as follows. If f is a polynomial time reduction from A to B then f is also a polynomial time reduction from \overline{A} to \overline{B} . And if f is a polynomial time reduction from A to B and g is a polynomial time reduction from B to C , then $g \circ f$ is a polynomial time reduction from A to C . \square

Finally, polynomial time reductions do serve as a way to compare the computational difficulty of two problems. We show that if $A \leq_P B$ and B is “easy” then A is easy.

Theorem A.47 *If $A \leq_P B$ and $B \in \mathbf{P}$ then $A \in \mathbf{P}$.*

Proof Let f be a polynomial time reduction from A to B and let M be a deterministic polynomial time algorithm recognizing B . Then the polynomial time algorithm N for A does the following: On input w , compute $f(w)$ and then run M on $f(w)$. It is easy to see that N recognizing A from the properties of a reduction.

The tricky step is to argue that N runs in polynomial time. Let us assume that f is computed in time n^k and let M run in time n^ℓ . Since f can be computed in time n^k , it means that $|f(w)| \leq |w|^k$; this is because a single step in the computation of f can produce at most one bit of $f(w)$. Therefore, the total running time of N is $|w|^k$ (time to compute $f(w)$) + $(|w|^k)^\ell$ (time to run M on $f(w)$ which is a string of length $|w|^k$). This is bounded by $O(n^{k\ell})$ which is polynomial. \square

In Theorem A.47, we could have replaced \mathbf{P} by any of the complexity classes in Fig. A.3 that contain \mathbf{P} , and the proof would go through. Thus, polynomial time reductions are an appropriate lens by which measure the relative difficulty of problems that belong to complexity classes that contain \mathbf{P} .

Definition A.48 (Hardness and Completeness)

Let C be a complexity class in Fig. A.3 that contains \mathbf{P} . A is said to be C -hard iff for every $B \in C$, $B \leq_P A$.

A is said to be C -complete iff $A \in C$ and A is C -hard.

In other words, informally, a problem is C -hard if it is at least as difficult as any problem in C . It is C -complete if in addition it also belongs to C . Fixing C to be NP , we could say that a problem is NP -complete if it is the “hardest” problem that belongs to NP . Because of their status as the most difficult problems in NP , they are candidate problems to study to help resolve the P versus NP question. This is captured by the following observation.

Proposition A.49 *If A is NP -hard and $A \in \text{P}$ then $\text{NP} = \text{P}$.*

Proof Consider any problem $B \in \text{NP}$. Since $B \leq_{\text{P}} A$ and $A \in \text{polytime}$, by Theorem A.47, we have $B \in \text{P}$. \square

In the absence of a firm resolution of the P versus NP questions, classifying a problem as NP -hard suggests that it is unlikely that the problem has a polynomial time algorithm given our belief that $\text{P} \neq \text{NP}$.

Many natural problems are NP -complete. The historically (and pedagogically) first problem known to be NP -complete is SAT (Cook-Levin Theorem Theorem 1.23).

Observe that from Fig. A.3, we have $\text{P} \subseteq \text{NP}$ and $\text{P} \subseteq \text{coNP}$. From this we can conclude that $\text{P} \subseteq \text{NP} \cap \text{coNP}$. Related but independent of the P versus NP question is whether $\text{P} = \text{NP} \cap \text{coNP}$. This question also remains open. Many problems that were previously known to be in $\text{NP} \cap \text{coNP}$ were proved to be in P years later. Two classical examples are Linear programming that was shown to be in P by Khachiyan in 1979 and testing whether a number is prime, which was proved by Agarwal-Kayal-Saxena in 2002 to be in P . However, there are some natural problems in $\text{NP} \cap \text{coNP}$ whose status with respect to P is still unresolved. One is the problem of solving parity games, and the other is the factoring problem.