

HEINZ-DIETER EBBINGHAUS
JÖRG FLUM

Finite Model Theory

Second Edition



Springer

Springer Monographs in Mathematics

Springer Monographs in Mathematics

Heinz-Dieter Ebbinghaus
Jörg Flum

Finite Model Theory

Second Revised and Enlarged
Edition 1999

 Springer

Heinz-Dieter Ebbinghaus

Jörg Flum

Mathematisches Institut

Abteilung für Mathematische Logik

Universität Freiburg

Eckerstraße 1

79104 Freiburg, Germany

e-mail: hde@math.uni-freiburg.de

joerg.flum@math.uni-freiburg.de

The softcover edition was published in 1999 by Springer-Verlag under the same title in the series *Perspectives in Mathematical Logic*.

Library of Congress Control Number: 2005932862

Mathematics Subject Classification (2000): 03C13, 03C80, 03D15, 68P15

ISSN 1439-7382

ISBN-10 3-540-28787-6 Springer Berlin Heidelberg New York

ISBN-13 978-3-540-28787-2 Springer Berlin Heidelberg New York

ISBN 3-540-65758-4 2nd Edition Springer-Verlag Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media

springeronline.com

© Springer-Verlag Berlin Heidelberg 1995, 1999, 2006

Printed in Germany

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

Typeset by the authors using a Springer \LaTeX macro package

Production: LE- \TeX Jelonek, Schmidt & Vöckler GbR, Leipzig

Cover design: Erich Kirchner, Heidelberg

Printed on acid-free paper

46/3142YL - 5 4 3 2 1 0

Preface

Finite model theory, the model theory of finite structures, has roots in classical model theory; however, its systematic development was strongly influenced by research and questions of complexity theory and of database theory.

Model theory or the theory of models, as it was first named by Tarski in 1954, may be considered as the part of the semantics of formalized languages that is concerned with the interplay between the syntactic structure of an axiom system on the one hand and (algebraic, settheoretic, ...) properties of its models on the other hand. As it turned out, first-order language (we mostly speak of first-order logic) became the most prominent language in this respect, the reason being that it obeys some fundamental principles such as the compactness theorem and the completeness theorem. These principles are valuable modeltheoretic tools and, at the same time, reflect the expressive weakness of first-order logic. This weakness is the breeding ground for the freedom which modeltheoretic methods rest upon.

By compactness, any first-order axiom system either has only finite models of limited cardinality or has infinite models. The first case is trivial because finitely many finite structures can explicitly be described by a first-order sentence. As model theory usually considers *all* models of an axiom system, modeltheorists were thus led to the second case, that is, to infinite structures. In fact, classical model theory of first-order logic and its generalizations to stronger languages live in the realm of the infinite. Basic methods such as the methods of constructing models, and basic aims such as a structure theory for first-order axiomatizable classes of structures are essentially concerned with infinite structures and transfinite or settheoretic combinatorics.

Nevertheless, there are natural reasons to consider finite structures. Historically, the most important one is the finite model property for certain classes of first-order formulas, that is, the equivalence of satisfiability and satisfiability in the finite. It was this property that settled the positive cases of the decision problem for prefix classes of first-order logic.

However, it took some twenty years until the middle of the twentieth century to really ask questions of a modeltheoretic flavour within the world of finite structures. The first landmarks here are Trahtenbrot's Theorem (1950) on the failure of the completeness theorem in the finite and the formulation

of the spectrum problem by Scholz (1952) that asks for a characterization of spectra in the finite of first-order sentences.

In both cases, computational aspects play their part: Trahtenbrot's proof rests on the undecidability of the halting problem for machines, and the spectrum problem turned out to be intimately linked to the question whether deterministic and nondeterministic polynomial time complexity coincide. The influence of computational aspects can be explained in various ways. First of all, finite structures can be coded as words and hence, can be objects of computations. Moreover, finite structures can serve to describe finite runs of machines. Finally, formulas of a logical language often can be interpreted as programs that, given a structure as input, perform the corresponding evaluation. This viewpoint is of importance, for example, in database theory, where relational databases are considered as finite structures.

We now give a short description of the contents of the book, at the same time motivating the choice of material and the emphasis it has been given.

Chapter 1 provides basic material concerning first-order logic. Chapters 2 and 3 are of a purely modeltheoretic character. When restricting oneself to finite structures, the essential theorems of first-order logic fail (this is documented at the end of Chapter 3) and important methods get trivial or useless. However, the gametheoretic methods of Ehrenfeucht and Fraïssé survive or even gain a special power. They are developed for first-order logic in Chapter 2 and for second-order and infinitary extensions in Chapter 3. Our representation is strongly based on isomorphism types (or, Hintikka formulas) as a unifying feature.

When turning to the finite, settheoretic combinatorics get replaced by finite combinatorics. As a consequence, there are new questions, for instance, questions of the kind: What is the relative frequency of graphs versus structures or of connected graphs versus graphs? Chapter 4 is dedicated to results aiming in this direction, so-called 0-1 laws, for first-order logic and some extensions. They say that for relational formulas either almost all or almost no finite structure is a model.

Chapter 5 treats the finite model property for some fragments of first-order logic, namely for the two-variable logic that consists of first-order formulas with only two variables, and for the class of $\forall\exists$ -formulas. The choice of the first logic is motivated by the methodological role of logics with a restricted number of variables, the choice of the second one by its relationship to concepts of the preceding chapter.

As has already been remarked, many questions in finite model theory are related to or even arose from questions in complexity theory. Chapter 6 gives an account of these connections on the computational level of finite automata.

When considering automata or even more powerful machines, it quickly becomes clear that first-order logic does not provide an adequate framework on the logical side because first-order logic lacks the ability of adequately expressing recursive procedures. This weakness of expressive power corre-

sponds to a similar phenomenon in classical model theory, and in a similar way as there, one can try to overcome this deficiency by introducing more powerful languages. In finite model theory, various so-called fixed-point logics have turned out to be promising candidates. They allow to speak about fixed-points of definable iteration procedures; the weakest one just allows to speak about the transitive closure of a definable binary relation.

Chapter 7 (together with Chapter 8) forms the core of the book. After a short introduction to fixed-point logics it develops what is known as the theory of descriptive complexity. Given a resource-bounded machine model, say the model of polynomially time-bounded Turing machines, and a suitable logic \mathcal{L} , say least fixed-point logic, then for every machine M of this kind there is a sentence φ of \mathcal{L} whose models are just the structures accepted by M , and for every sentence φ of \mathcal{L} there is a machine M of this kind that just accepts the models of φ . Hence, the classes of structures that are acceptable by a polynomially time-bounded machine, correspond to the classes that are axiomatizable in \mathcal{L} . In this way, one obtains logical descriptions of complexity classes and, therefore, logical analogues of major problems in complexity theory. For example, the $\text{PTIME} = \text{PSPACE}$ -problem now amounts to the question whether two fixed-point logics have the same expressive power in finite structures. The chapter gives an account of the most important results of this kind. As it turns out, they mostly presuppose that the structures considered are ordered. This assumption seems to be natural, since we induce an ordering on the universe of a structure, when encoding it as an input string. On the other hand, some major open problems of descriptive complexity theory are concerned with the question to what extent the descriptive results can be generalized to structures without an ordering.

The proof indicated above, for the fact that least fixed-point logic corresponds to PTIME , yields that any sentence φ of this logic is equivalent to a sentence describing the behaviour of a machine and hence, has a special syntactical structure. So, as a byproduct of the logical characterizations of complexity classes, one gets certain normal forms for the logics involved.

Chapter 8 (with parts of Chapter 9) presents the model theory of fixed-point logics for finite structures. The material presented includes a thorough discussion of transitive closure logics, as well as connections to second-order logic and infinitary logics.

Chapter 9 takes up the viewpoint of formulas as programs or queries, introducing a bunch of programming languages of the DATALOG family in database theory. The main aim here is to show the equivalence of these languages with certain fixed-point languages, thus opening another methodological gateway to the latter ones. The way pays, as can be seen by new and far reaching possibilities of obtaining normal forms. The chapter concludes with the investigation of the fine structure of fixed-point languages.

Chapter 10 lies outside the mainstream of the book. It is concerned with a logically-oriented representation of optimization problems, but is restricted to basic material.

Chapter 11 is concerned with one of the most prominent open problems of finite model theory, namely whether there is a logic that captures PTIME also on unordered structures. As it turns out, this question is linked to the problem of feasibly characterizing finite structures by invariants. By coding structures in graphs we reduce the problem to the case of graphs. The chapter ends by providing positive answers for some classes of graphs, thus illustrating an at the present time active direction of research.

Finally, Chapter 12 discusses a concept that is well-known from classical model theory, the concept of a quantifier. The idea here is the following: In order to find a logic \mathcal{L} that corresponds to a complexity class \mathcal{C} in the sense described above, one can try to start with a simple logic, say first-order logic, and add a quantifier that incorporates a \mathcal{C} -complete problem. When pursuing this idea, one has to find analogues of notions from complexity theory such as reductions, hardness, completeness, etc. The chapter concludes with an example of the analogy between quantifiers in logic and oracles in complexity theory.

As the preceding description shows, the core of the book is centered around modeltheoretic issues related to descriptive complexity. Such a concentration seemed to be necessary in order to come along with a book of a reasonable size. Among the major gaps the reader will encounter we mention two ones lying at opposite ends of the spectrum of possible topics: the theory of circuits on the more computational side and the work related to a structure theory for the finite on the more modeltheoretic side. Concerning topics of the first group we refer the reader to Neil Immerman's recent monograph "Computational Complexity".

Some chapters are as independent from each other as their contents allows. In particular, a reader interested in descriptive complexity theory should not have problems to start with Chapter 7. And a modeltheorist interested in the finite model theory of fixed-point logics can start with Chapter 8.

This second edition is a thoroughly revised and enlarged version of the original text, the most relevant addition being the new Chapter 11. Moreover, we have tried to take into consideration criticism and suggestions of readers of the first edition. Again, with gratitude we mention Martin Grohe: his proposals and his results have influenced and enriched considerable parts of the text.

Table of Contents

Preface	V
1. Preliminaries	1
2. The Ehrenfeucht-Fraïssé Method	13
2.1 Elementary Classes	13
2.2 Ehrenfeucht's Theorem	15
2.3 Examples and Fraïssé's Theorem	20
2.4 Hanf's Theorem	26
2.5 Gaifman's Theorem	30
3. More on Games	37
3.1 Second-Order Logic	37
3.2 Infinitary Logic: The Logics $L_{\infty\omega}$ and $L_{\omega_1\omega}$	40
3.3 The Logics FO^s and $L_{\infty\omega}^s$	46
3.3.1 Pebble Games	49
3.3.2 The s -Invariant of a Structure	54
3.3.3 Scott Formulas	56
3.4 Logics with Counting Quantifiers	58
3.5 Failure of Classical Theorems in the Finite	62
4. 0–1 Laws	71
4.1 0–1 Laws for FO and $L_{\infty\omega}^\omega$	71
4.2 Parametric Classes	74
4.3 Unlabeled 0–1 Laws	77
4.3.1 Appendix	82
4.4 Examples and Consequences	84
4.5 Probabilities of Monadic Second Order Properties	88
5. Satisfiability in the Finite	95
5.1 Finite Model Property of FO^2	95
5.2 Finite Model Property of $\forall^2\exists^*$ -Sentences	99

6. Finite Automata and Logic: A Microcosm of Finite Model Theory	105
6.1 Languages Accepted by Automata	105
6.2 Word Models	108
6.3 Examples and Applications	111
6.4 First-Order Definability	114
7. Descriptive Complexity Theory	119
7.1 Some Extensions of First-Order Logic	120
7.2 Turing Machines and Complexity Classes	124
7.2.1 Digression: Trahtenbrot's Theorem	127
7.2.2 Structures as Inputs	129
7.3 Logical Descriptions of Computations	133
7.4 The Complexity of the Satisfaction Relation	147
7.5 The Main Theorem and Some Consequences	151
7.5.1 Appendix	162
8. Logics with Fixed-Point Operators	165
8.1 Inflationary and Least Fixed-Points	165
8.2 Simultaneous Induction and Transitivity	177
8.3 Partial Fixed-Point Logic	191
8.4 Fixed-Point Logics and $L_{\infty\omega}^{\omega}$	198
8.4.1 The Logic $FO(PFP_{PTIME})$	205
8.4.2 Fixed-Point Logic with Counting	207
8.5 Fixed-Point Logics and Second-Order Logic	210
8.5.1 Digression: Implicit Definability	217
8.6 Transitive Closure Logic	220
8.6.1 $FO(DTC) < FO(TC)$	221
8.6.2 $FO(posTC)$ and Normal Forms	224
8.6.3 $FO(TC) < FO(LFP)$	229
8.7 Bounded Fixed-Point Logic	235
9. Logic Programs	239
9.1 DATALOG	239
9.2 I-DATALOG and P-DATALOG	245
9.3 A Preservation Theorem	250
9.4 Normal Forms for Fixed-Point Logics	253
9.5 An Application of Negative Fixed-Point Logic	263
9.6 Hierarchies of Fixed-Point Logics	268
10. Optimization Problems	275
10.1 Polynomially Bounded Optimization Problems	275
10.2 Approximable Optimization Problems	280

11. Logics for PTIME	287
11.1 Logics and Invariants	288
11.2 PTIME on Classes of Structures	295
12. Quantifiers and Logical Reductions	307
12.1 Lindström Quantifiers	308
12.2 PTIME and Quantifiers	314
12.3 Logical Reductions	320
12.4 Quantifiers and Oracles	330
References	339
Index	349

1. Preliminaries

The purpose of this section is to fix notations and terminology for the basic notions related to first-order logic. We assume that the reader has already some familiarity with these concepts, as it is obtained by an introductory course in mathematical logic. For more detailed information we refer to textbooks such as [32, 35]. Parts of our exposition in part B follow [9].

A Structures

Vocabularies are finite sets that consist of *relation symbols* P, Q, R, \dots and *constant symbols* (for short: *constants*) c, d, \dots ¹ Every relation symbol is equipped with a natural number ≥ 1 , its *arity*. We denote vocabularies by τ, σ, \dots . A vocabulary is *relational*, if it does not contain constants.

A *structure* \mathcal{A} of vocabulary τ (by short: a τ -*structure*) consists of a nonempty set A , the *universe* or *domain* of \mathcal{A} , of an n -ary relation $R^{\mathcal{A}}$ on A for every n -ary relation symbol R in τ , and of an element $c^{\mathcal{A}}$ of A for every constant c in τ . (Mostly we use the notations $R^{\mathcal{A}}$ for $R^{\mathcal{A}}$ and $c^{\mathcal{A}}$ for $c^{\mathcal{A}}$.) An n -ary *relation* $R^{\mathcal{A}}$ on A is a subset of A^n , the set of n -tuples of elements of A . We mostly write $R^{\mathcal{A}}a_1 \dots a_n$ instead of $(a_1, \dots, a_n) \in R^{\mathcal{A}}$. A structure \mathcal{A} is *finite*, if its universe A is a finite set.

We give some examples of structures that will play a prominent role in the book.

A1 Graphs


Let $\tau = \{E\}$ with a binary relation symbol E . A *graph* (or, *undirected graph*) is a τ -structure $\mathcal{G} = (G, E^{\mathcal{G}})$ satisfying

- (1) for all $a \in G$: not $E^{\mathcal{G}}aa$
- (2) for all $a, b \in G$: if $E^{\mathcal{G}}ab$ then $E^{\mathcal{G}}ba$.

By GRAPH we denote the class of *finite* graphs. If only (1) is required, we speak of a *digraph* (or, *directed graph*). The elements of G are sometimes called *points* or *vertices*, the elements of $E^{\mathcal{G}}$ *edges*. The following figures represent

¹ Usually also function symbols are allowed in vocabularies. For the purposes of this book our definition does not represent an essential restriction. We explain in part D how the results can be extended to vocabularies with function symbols.

the graph $(\{a, b, c, d\}, \{(a, b), (b, a), (b, c), (c, b), (b, d), (d, b), (c, d), (d, c)\})$ and the digraph $(\{a, b, c, d\}, \{(a, b), (b, a), (b, c), (b, d), (d, c)\})$.



A subset X of the universe of a graph \mathcal{G} is a *clique*, if $E^{\mathcal{G}}ab$ for all $a, b \in X$, $a \neq b$.

Let \mathcal{G} be a digraph. If $n \geq 1$ and

$$E^{\mathcal{G}}a_0a_1, E^{\mathcal{G}}a_1a_2, \dots, E^{\mathcal{G}}a_{n-1}a_n$$

then a_0, \dots, a_n is a *path* from a_0 to a_n of *length* n . If $a_0 = a_n$ then a_0, \dots, a_n is a *cycle*. \mathcal{G} is *acyclic* if it has no cycle. A path a_0, \dots, a_n is *Hamiltonian* if $G = \{a_0, \dots, a_n\}$ and $a_i \neq a_j$ for $i \neq j$. If, in addition, $E^{\mathcal{G}}a_n a_0$ we speak of a *Hamiltonian circuit*.

Let \mathcal{G} be a graph. Write $a \sim b$ if $a = b$ or if there is a path from a to b . Clearly, \sim is an equivalence relation. The equivalence class of a is called the (*connected*) *component* of a . \mathcal{G} is *connected* if $a \sim b$ for all $a, b \in G$, that is, if there is only one connected component. Let CONN be the class of finite connected graphs.

Denote by $d(a, b)$ the length of a shortest path from a to b ; more precisely, define the *distance function* $d : G \times G \rightarrow \mathbb{N} \cup \{\infty\}$ by ²

$$d(a, b) = \infty \text{ iff } a \not\sim b; \quad d(a, b) = 0 \text{ iff } a = b;$$

and otherwise,

$$d(a, b) = \min\{n \geq 1 \mid \text{there is a path from } a \text{ to } b \text{ of length } n\}.$$

Obviously,

$$d(a, c) \leq d(a, b) + d(b, c),$$

where we use the natural conventions for ∞ .

For simplicity, we give the following definitions only for *finite* digraphs. A vertex b is a *successor* of a vertex a (and a a *predecessor* of b) if $E^{\mathcal{G}}ab$. The *in-degree* of a vertex is the number of its predecessors, the *out-degree* the number of its successors. In graphs the in-degree and the out-degree of a vertex a coincide and are called the *degree* of a .

² \mathbb{N} denotes the set of natural numbers, $\mathbb{N} = \{0, 1, 2, \dots\}$.

A *root* of a digraph is a vertex with in-degree 0 and a *leaf* a vertex with out-degree 0.

A *forest* is an acyclic digraph where each vertex has in-degree at most 1. A *tree* is a forest with connected underlying graph, i.e., a forest (G, E^G) such that $(G, \{(a, b) \mid E^G ab \text{ or } E^G ba\})$ is connected. Note that a finite tree has exactly one root. Let TREE be the class of finite trees.

A2 Orderings

Let $\tau = \{<\}$ with a binary relation symbol $<$. A τ -structure $\mathcal{A} = (A, <^A)$ is called an *ordering* if for all $a, b, c \in A$:

- (1) not $a <^A a$
- (2) $a <^A b$ or $b <^A a$ or $a = b$
- (3) if $a <^A b$ and $b <^A c$ then $a <^A c$.

Sometimes we consider finite orderings also as $\{<, S, \min, \max\}$ -structures. Here S is a binary relation symbol representing the successor relation and \min and \max are constants for the first and the last element of the ordering. (When considering the natural ordering on $\{0, \dots, n\}$ we often refer to \min as the zero-th element.) Thus, a finite $\{<, S, \min, \max\}$ -structure \mathcal{A} is an ordering if, in addition to (1), (2), (3), for all $a, b \in A$:

- (4) $S^A ab$ iff ($a <^A b$ and for all c , if $a <^A c$ then $b <^A c$ or $b = c$)
- (5) $\min^A <^A a$ or $\min^A = a$
- (6) $a <^A \max^A$ or $a = \max^A$.

For other purposes it might be advantageous to consider finite orderings as $\{<, \min, \max\}$ -structures. Suppose that τ_0 is a vocabulary with $\{<\} \subseteq \tau_0 \subseteq \{<, S, \min, \max\}$ and let σ be an arbitrary vocabulary with $\tau_0 \subseteq \sigma$. Later in this book the class $\mathcal{O}[\sigma]$ of finite *ordered* σ -structures will play a prominent role. Here a finite σ -structure \mathcal{A} is said to be *ordered*, if the *reduct* $\mathcal{A}|_{\tau_0}$ (i.e., the τ_0 -structure obtained from \mathcal{A} by forgetting the interpretations of the symbols in $\sigma \setminus \tau_0$) is an ordering.

A3 Operations on Structures

Two τ -structures \mathcal{A} and \mathcal{B} are *isomorphic*, written $\mathcal{A} \cong \mathcal{B}$, if there is an *isomorphism* from \mathcal{A} to \mathcal{B} , i.e., a bijection $\pi: A \rightarrow B$ preserving relations and constants, that is,

- for n -ary $R \in \tau$ and $a_1, \dots, a_n \in A$,

$$R^A a_1 \dots a_n \quad \text{iff} \quad R^B \pi(a_1) \dots \pi(a_n)$$

- for $c \in \tau$, $\pi(c^A) = c^B$.

For τ -structures \mathcal{A} and \mathcal{B} , the *product* $\mathcal{A} \times \mathcal{B}$ of \mathcal{A} and \mathcal{B} is the τ -structure with domain $A \times B := \{(a, b) \mid a \in A, b \in B\}$, which is given by

– for n -ary R in τ and $(a_1, b_1), \dots, (a_n, b_n) \in A \times B$,

$$R^{A \times B}(a_1, b_1) \dots (a_n, b_n) \quad \text{iff} \quad R^A a_1 \dots a_n \text{ and } R^B b_1 \dots b_n$$

– for c in τ , $c^{A \times B} := (c^A, c^B)$.

For relational τ we introduce the *union* (or, *disjoint union*) of structures. Assume that \mathcal{A} and \mathcal{B} are τ -structures with $A \cap B = \emptyset$. Then $\mathcal{A} \dot{\cup} \mathcal{B}$, the *union* of \mathcal{A} and \mathcal{B} , is the τ -structure with domain $A \cup B$ and

$$R^{\mathcal{A} \dot{\cup} \mathcal{B}} := R^{\mathcal{A}} \cup R^{\mathcal{B}}$$

for any R in τ . In case \mathcal{A} and \mathcal{B} are structures with $A \cap B \neq \emptyset$, we take isomorphic copies \mathcal{A}' of \mathcal{A} and \mathcal{B}' of \mathcal{B} with disjoint universes (e.g., with universes $A \times \{1\}$ and $B \times \{2\}$) and set $\mathcal{A} \dot{\cup} \mathcal{B} := \mathcal{A}' \dot{\cup} \mathcal{B}'$.

Note that the union of ordered structures is not an ordered structure. The situation is different for the so-called *ordered sum*: Let τ with $< \in \tau$ be relational and let \mathcal{A} and \mathcal{B} be ordered τ -structures. Assume that $A \cap B = \emptyset$. Define $\mathcal{A} \triangleleft \mathcal{B}$, the *ordered sum* of \mathcal{A} and \mathcal{B} , as $\mathcal{A} \dot{\cup} \mathcal{B}$ but setting

$$<^{\mathcal{A} \triangleleft \mathcal{B}} := <^{\mathcal{A}} \cup <^{\mathcal{B}} \cup \{(a, b) \mid a \in A, b \in B\},$$

that is, in $\mathcal{A} \triangleleft \mathcal{B}$ all elements of A precede all elements of B .

Note that $\mathcal{A} \times \mathcal{B} \cong \mathcal{B} \times \mathcal{A}$, $\mathcal{A} \dot{\cup} \mathcal{B} \cong \mathcal{B} \dot{\cup} \mathcal{A}$ but, in general, $\mathcal{A} \triangleleft \mathcal{B} \not\cong \mathcal{B} \triangleleft \mathcal{A}$. The definition of product, union, and ordered sum can easily be extended to more than two structures setting, for example, $\mathcal{A} \triangleleft \mathcal{B} \triangleleft \mathcal{C} := ((\mathcal{A} \triangleleft \mathcal{B}) \triangleleft \mathcal{C})$. For a finite nonempty set I we denote by \mathcal{A}^I , $\dot{\cup}_I \mathcal{A}$, and $\triangleleft_I \mathcal{A}$ the product, the union, and the ordered sum, respectively, of $\|I\|$ copies of \mathcal{A} . Here $\|I\|$ denotes the cardinality of I .

B Syntax and Semantics of First-Order Logic

We now turn to the syntactic notions of first-order logic FO. Fix a vocabulary τ . Each formula of first-order logic will be a string of symbols taken from the alphabet consisting of

- v_1, v_2, v_3, \dots (the *variables*)
- \neg, \vee (the *connectives not, or*)
- \exists (the *existential quantifier*)
- $=$ (the *equality symbol*)
- $), ($
- the symbols in τ .

A *term* of vocabulary τ is a variable or a constant in τ . Henceforth, we often shall use the letters x, y, z, \dots for variables and t, t_1, \dots for terms.

The *formulas* of first-order logic of vocabulary τ are those strings which are obtained by finitely many applications of the following rules:

- (F1) If t_0 and t_1 are terms then $t_0 = t_1$ is a formula.
(F2) If R in τ is n -ary and t_1, \dots, t_n are terms then $Rt_1 \dots t_n$ is a formula.
(F3) If φ is a formula then $\neg\varphi$ is a formula.
(F4) If φ and ψ are formulas then $(\varphi \vee \psi)$ is a formula.
(F5) If φ is a formula and x a variable then $\exists x\varphi$ is a formula.

Denote by $\text{FO}[\tau]$ the set of formulas of first-order logic of vocabulary τ . Formulas obtained by (F1) or (F2) are called *atomic* formulas. For formulas φ and ψ we use $(\varphi \wedge \psi)$, $(\varphi \rightarrow \psi)$, $(\varphi \leftrightarrow \psi)$, and $\forall x\varphi$ as abbreviations for the formulas $\neg(\neg\varphi \vee \neg\psi)$, $(\neg\varphi \vee \psi)$, $((\neg\varphi \vee \psi) \wedge (\neg\psi \vee \varphi))$, and $\neg\exists x\neg\varphi$, respectively.

We shall often omit parentheses in formulas if they are not essential such as the outermost parentheses in disjunctions $(\varphi \vee \psi)$. In examples, different letters $x, y, z \dots$ will always stand for different variables.

The axioms for graphs stated above have the following formalizations in $\text{FO}[\{E\}]$:

$$\begin{aligned} &\forall x\neg Exx \\ &\forall x\forall y(Exy \rightarrow Eyx), \end{aligned}$$

those for orderings the following formalizations in $\text{FO}[\{<\}]$:

$$\begin{aligned} &\forall x\neg x < x \\ &\forall x\forall y(x < y \vee y < x \vee x = y) \\ &\forall x\forall y\forall z((x < y \wedge y < z) \rightarrow x < z). \end{aligned}$$

For orderings as $\{<, S, \min, \max\}$ -structures we need in addition

$$\begin{aligned} &\forall x\forall y(Sxy \leftrightarrow (x < y \wedge \forall z(x < z \rightarrow (y < z \vee y = z)))) \\ &\forall x(\min < x \vee \min = x) \\ &\forall x(x < \max \vee x = \max). \end{aligned}$$

All these formulas are *sentences*, i.e., formulas in which every variable in an atomic subformula is in the scope of a corresponding quantifier. Such occurrences are called *bound* occurrences. The last occurrence of x in

$$(\forall x\neg Exx \wedge \exists yEyx)$$

is not in the scope of a quantifier binding it. Such occurrences are called *free*. The notion of a *free* variable of a formula φ is made precise by the following definition, a definition by *induction on* (the length of) φ : The set $\text{free}(\varphi)$ of *free variables* of a formula φ is defined by:

- If φ is atomic then the set $\text{free}(\varphi)$ of free variables of φ is the set of variables occurring in φ
- $\text{free}(\neg\varphi) := \text{free}(\varphi)$
- $\text{free}(\varphi \vee \psi) := \text{free}(\varphi) \cup \text{free}(\psi)$
- $\text{free}(\exists x\varphi) := \text{free}(\varphi) \setminus \{x\}$.

It is common practice to use the notation $\varphi(x_1, \dots, x_n)$ to indicate that x_1, \dots, x_n are distinct and $\text{free}(\varphi) \subseteq \{x_1, \dots, x_n\}$ without implying that all x_i are actually free in φ (in some chapters of the book we shall give another meaning to this notation). Often we abbreviate an n -tuple x_1, \dots, x_n of variables by \bar{x} , for example, writing $\varphi(\bar{x})$ for $\varphi(x_1, \dots, x_n)$. Usually we do not make explicit the length of \bar{x} (here n), its size either being inessential or clear from the context. Moreover, we often omit commas writing, for example, $\bar{x} = x_1 \dots x_n$.

So far, terms and formulas of FO are simply finite strings of symbols. We now assign the intended meanings to the logical symbols so that, in particular, the sentences above really formalize the axioms of graphs and orderings. We do this by defining the *satisfaction relation* $\mathcal{A} \models \varphi$ between structures on the one hand and sentences on the other hand.

Let \mathcal{A} be a τ -structure. An *assignment* in \mathcal{A} is a function α with domain the set of variables and with values in A , $\alpha : \{v_n \mid n \geq 1\} \rightarrow A$. Think of α as assigning the meaning $\alpha(x)$ to the variable x . Extend α to a function defined for all terms by setting $\alpha(c) := c^{\mathcal{A}}$ for c in τ . Denote by $\alpha \frac{a}{x}$ the assignment that agrees with α except that $\alpha \frac{a}{x}(x) = a$.

We define the relation

$$\mathcal{A} \models \varphi[\alpha]$$

(“the assignment α *satisfies* the formula φ in \mathcal{A} ” or “ φ is *true* in \mathcal{A} under α ”) as follows:

$$\begin{array}{lll} \mathcal{A} \models t_1 = t_2[\alpha] & \text{iff} & \alpha(t_1) = \alpha(t_2) \\ \mathcal{A} \models R t_1 \dots t_n[\alpha] & \text{iff} & R^{\mathcal{A}} \alpha(t_1) \dots \alpha(t_n) \\ \mathcal{A} \models \neg \varphi[\alpha] & \text{iff} & \text{not } \mathcal{A} \models \varphi[\alpha] \\ \mathcal{A} \models (\varphi \vee \psi)[\alpha] & \text{iff} & \mathcal{A} \models \varphi[\alpha] \text{ or } \mathcal{A} \models \psi[\alpha] \\ \mathcal{A} \models \exists x \varphi[\alpha] & \text{iff} & \text{there is an } a \in A \text{ such that } \mathcal{A} \models \varphi[\alpha \frac{a}{x}]. \end{array}$$

Note that the truth or falsity of $\mathcal{A} \models \varphi[\alpha]$ depends only on the values of α for those variables x which are free in φ . That is, if $\alpha_1(x) = \alpha_2(x)$ for all $x \in \text{free}(\varphi)$, then $\mathcal{A} \models \varphi[\alpha_1]$ iff $\mathcal{A} \models \varphi[\alpha_2]$. Thus, if $\varphi = \varphi(x_1, \dots, x_n)$ and $a_1 = \alpha(x_1), \dots, a_n = \alpha(x_n)$, then we may write $\mathcal{A} \models \varphi[a_1, \dots, a_n]$ for $\mathcal{A} \models \varphi[\alpha]$. In particular, if φ is a sentence, then the truth or falsity of $\mathcal{A} \models \varphi[\alpha]$ is completely independent of α . Thus we may write $\mathcal{A} \models \varphi$ (read: \mathcal{A} is a *model* of φ , or \mathcal{A} *satisfies* φ), if for some (hence every) assignment α , $\mathcal{A} \models \varphi[\alpha]$. For a set Φ of formulas, $\mathcal{A} \models \Phi[\alpha]$ means that $\mathcal{A} \models \varphi[\alpha]$ for all $\varphi \in \Phi$. Φ is *satisfiable* if there is a structure \mathcal{A} and an assignment α in \mathcal{A} such that $\mathcal{A} \models \Phi[\alpha]$.

A formula ψ is a *consequence* of Φ , written $\Phi \models \psi$, if $\mathcal{A} \models \psi[\alpha]$ whenever $\mathcal{A} \models \Phi[\alpha]$. The formula ψ is *logically valid*, written $\models \psi$, if $\emptyset \models \psi$, that is, if ψ is true in all structures under all assignments. And formulas φ and ψ are *logically equivalent* if $\models \varphi \leftrightarrow \psi$. When only taking into consideration finite

structures, we use the notations $\Phi \models_{\text{fin}} \psi$ and $\models_{\text{fin}} \psi$, and speak of *equivalent* formulas; hence, φ and ψ are equivalent if $\models_{\text{fin}} \varphi \leftrightarrow \psi$, that is, if $\varphi \leftrightarrow \psi$ holds in all *finite* structures under all assignments.

At some places it will be convenient to assume that first-order logic contains two zero-ary relation symbols T, F. In every structure, T and F are interpreted as TRUE (i.e., as being true) and FALSE, respectively. Hence, the atomic formula T is logically equivalent to $\exists xx = x$ and F to $\neg \exists xx = x$. If $\Phi = \{\varphi_1, \dots, \varphi_n\}$ we sometimes write $\bigwedge \Phi$ for $\varphi_1 \wedge \dots \wedge \varphi_n$ and $\bigvee \Phi$ for $\varphi_1 \vee \dots \vee \varphi_n$.³ In case $\Phi = \emptyset$ we set $\bigwedge \Phi = \text{T}$ and $\bigvee \Phi = \text{F}$. Then, for arbitrary finite Φ ,

$$\mathcal{A} \models \bigwedge \Phi \quad \text{iff} \quad \text{for all } \varphi \in \Phi, \mathcal{A} \models \varphi.$$

We introduce some further syntactic notions and notations. The *quantifier rank* $\text{qr}(\varphi)$ of a formula φ is the maximum number of nested quantifiers occurring in it:

$$\begin{aligned} \text{qr}(\varphi) &:= 0, & \text{if } \varphi \text{ is atomic;} & & \text{qr}(\neg\varphi) &:= \text{qr}(\varphi); \\ \text{qr}(\varphi \vee \psi) &:= \max\{\text{qr}(\varphi), \text{qr}(\psi)\}; & & & \text{qr}(\exists x\varphi) &:= \text{qr}(\varphi) + 1. \end{aligned}$$

It can be shown that every first-order formula is logically equivalent to a formula in prenex normal form, that is, to a formula of the form $Q_1x_1 \dots Q_sx_s\psi$, where $Q_1, \dots, Q_s \in \{\forall, \exists\}$ and where ψ is quantifier-free. Such a formula is called Σ_n , if the string of quantifiers consists of n consecutive blocks, where in each block all quantifiers are of the same type (i.e., all universal or all existential), adjacent blocks contain quantifiers of different type, and the first block is existential. Π_n formulas are defined in the same way, but now we require that the first block consists of universal quantifiers. A Δ_n -formula is a formula logically equivalent to both a Σ_n -formula and a Π_n -formula.

If $\varphi(x_1, \dots, x_n)$ is a formula and t_1, \dots, t_n are terms then $\varphi \frac{t_1 \dots t_n}{x_1 \dots x_n}$ or, more simply, $\varphi(t_1, \dots, t_n)$ denotes the result of simultaneously replacing all free occurrences of x_1, \dots, x_n by t_1, \dots, t_n , respectively. This presupposes that none of the variables in t_1, \dots, t_n gets into the scope of a corresponding quantifier; otherwise, the bound variables in φ must be renamed in some canonical fashion before replacing.

Given a formula $\varphi(x, \bar{z})$ and $n \geq 1$,

$$\exists^{\geq n} x \varphi(x, \bar{z})$$

is an abbreviation for the formula

$$\exists x_1 \dots \exists x_n \left(\bigwedge_{1 \leq i \leq n} \varphi(x_i, \bar{z}) \wedge \bigwedge_{1 \leq i < j \leq n} \neg x_i = x_j \right)$$

³ For definiteness, given the vocabulary τ , fix an ordering on the alphabet of first-order logic and interpret a conjunction $\bigwedge \Phi$ for a finite set Φ as the iterative conjunction of the elements of Φ in the induced lexicographic ordering, say.

expressing that there are at least n elements x with $\varphi(x, \bar{z})$. $\exists^=n x \varphi(x, \bar{z})$ and $\exists^{\leq n} x \varphi(x, \bar{z})$ are defined similarly. Moreover, we set

$$\varphi_{\geq n} := \exists^{\geq n} x x = x; \quad \varphi_{=n} := \exists^=n x x = x; \quad \varphi_{\leq n} := \exists^{\leq n} x x = x.$$

Clearly,

$$\mathcal{A} \models \varphi_{\geq n} \quad \text{iff} \quad \|\mathcal{A}\| \geq n$$

and, similarly, for $\varphi_{=n}$ and $\varphi_{\leq n}$.

C Some Classical Results of First-Order Logic

First-order logic has been used as a framework to analyze the notion of mathematical proof. A result of this analysis is Gödel's Completeness Theorem: A mathematical proof of ψ from Φ shows that $\Phi \models \psi$, i.e. that ψ is a consequence of Φ . A natural question is whether the converse holds, too, that is, whether for any consequence there is a proof. To give an answer, Gödel used a notion of *formal proof* that is based on a finite system of formal rules. A *formal proof* of ψ from Φ consists of a sequence of applications of these rules leading from the formulas in Φ to ψ .⁴ In a natural way every formal proof corresponds to a mathematical proof. Thus, if ψ is formally provable from Φ then ψ is a consequence of Φ . Moreover, Gödel showed:

Theorem 1.0.1 (Completeness Theorem) *ψ is a consequence of Φ iff ψ is formally provable from Φ .*

Two immediate consequences are:

Theorem 1.0.2 *The set of logically valid sentences of first-order logic is recursively enumerable.*

Theorem 1.0.3 (Compactness Theorem) (a) *If ψ is a consequence of Φ then ψ is already a consequence of a finite subset of Φ .*

(b) *If every finite subset of Φ is satisfiable then Φ is satisfiable.*

The proof of the Completeness Theorem often leads to a proof of

Theorem 1.0.4 (Löwenheim-Skolem Theorem) *If Φ has a model then Φ has an at most countable model.*

Neither 1.0.2 nor 1.0.3 remain valid if one only considers finite structures. A counterexample for the Compactness Theorem is given by the set $\Phi_{\infty} := \{\varphi_{\geq n} \mid n \geq 1\}$: Each finite subset of Φ_{∞} has a finite model, but Φ_{∞} has no finite model.

The failure of 1.0.2 is documented by

Theorem 1.0.5 (Traftenbrot's Theorem) *The set of sentences of first-order logic valid in all finite structures is not recursively enumerable.*

⁴ We do not give detailed definitions, since we do not need them later.

A proof will be given in section 7.2.

We derive a consequence of the Compactness Theorem, which will be used in section 2.5.

Lemma 1.0.6 *Let $\varphi \in \text{FO}[\tau]$ and for $i \in I$, let $\Phi^i \subseteq \text{FO}[\tau]$. Assume that*

$$(*) \quad \models \varphi \leftrightarrow \bigvee_{i \in I} \bigwedge \Phi^i.^5$$

Then there is a finite $I_0 \subseteq I$ and, for every $i \in I_0$, a finite $\Phi_0^i \subseteq \Phi^i$ such that

$$\models \varphi \leftrightarrow \bigvee_{i \in I_0} \bigwedge \Phi_0^i.$$

Proof. For simplicity we assume that φ is a sentence and that every Φ^i is a set of sentences. By hypothesis, for $i \in I$ we have $\Phi^i \models \varphi$; hence, by the Compactness Theorem, $\Phi_0^i \models \varphi$ for some finite $\Phi_0^i \subseteq \Phi^i$, and therefore, $\models \bigvee_{i \in I_0} \bigwedge \Phi_0^i \rightarrow \varphi$ for each finite subset $I_0 \subseteq I$. If there is no such I_0 with $\models \varphi \rightarrow \bigvee_{i \in I_0} \bigwedge \Phi_0^i$, then each finite subset of $\{\varphi\} \cup \{\neg \bigwedge \Phi_0^i \mid i \in I\}$ has a model. Hence, by the Compactness Theorem, there is a model of φ which for all $i \in I$ satisfies $\neg \bigwedge \Phi_0^i$. This contradicts (*). \square

Structures \mathcal{A} and \mathcal{B} (of the same vocabulary) are said to be *elementarily equivalent*, written $\mathcal{A} \equiv \mathcal{B}$, if they satisfy the same first-order sentences. The preceding lemma has the following

Corollary 1.0.7 *Let Φ be a set of first-order sentences. Assume that any two structures that satisfy the same sentences of Φ are elementarily equivalent. Then any first-order sentence is equivalent to a boolean combination of sentences of Φ (that is, is equivalent to a sentence obtainable by closing Φ under \neg and \vee).*

Proof. For any structure \mathcal{A} set

$$\Phi(\mathcal{A}) := \{\psi \mid \psi \in \Phi, \mathcal{A} \models \psi\} \cup \{\neg\psi \mid \psi \in \Phi, \mathcal{A} \models \neg\psi\}.$$

Let φ be any first-order sentence. By the preceding lemma it suffices to show that

$$\models \varphi \leftrightarrow \bigvee_{\mathcal{A} \models \varphi} \bigwedge \Phi(\mathcal{A}).$$

Clearly, if $\mathcal{B} \models \varphi$ then $\mathcal{B} \models \bigvee_{\mathcal{A} \models \varphi} \bigwedge \Phi(\mathcal{A})$, since $\mathcal{B} \models \Phi(\mathcal{B})$. For the converse, suppose $\mathcal{B} \models \bigvee_{\mathcal{A} \models \varphi} \bigwedge \Phi(\mathcal{A})$. Then for some model \mathcal{A} of φ , $\mathcal{B} \models \Phi(\mathcal{A})$. By definition of $\Phi(\mathcal{A})$, \mathcal{A} and \mathcal{B} satisfy the same sentences of Φ and hence, by hypothesis, are elementarily equivalent. Therefore, $\mathcal{B} \models \varphi$. \square

⁵ This means that under any assignment in any structure, φ is true iff for some $i \in I$ all formulas in Φ^i are true.

D Model Classes and Global Relations

Fix a vocabulary τ . For a sentence φ of $\text{FO}[\tau]$ we denote by $\text{Mod}(\varphi)$ the class of *finite* models of φ . If π is an isomorphism from \mathcal{A} to \mathcal{B} , $\varphi(x_1, \dots, x_n) \in \text{FO}[\tau]$, and $a_1, \dots, a_n \in A$, then an easy induction on formulas shows

$$(1) \quad \mathcal{A} \models \varphi[a_1, \dots, a_n] \quad \text{iff} \quad \mathcal{B} \models \varphi[\pi(a_1), \dots, \pi(a_n)].$$

In particular, if φ is a sentence then

$$\mathcal{A} \models \varphi \quad \text{iff} \quad \mathcal{B} \models \varphi.$$

Hence, $\text{Mod}(\varphi)$ is closed under isomorphisms, that is,

$$(2) \quad \mathcal{A} \in \text{Mod}(\varphi) \text{ and } \mathcal{A} \cong \mathcal{B} \text{ imply } \mathcal{B} \in \text{Mod}(\varphi).$$

For $\varphi(x_1, \dots, x_n) \in \text{FO}[\tau]$ and a structure \mathcal{A} let

$$\varphi^{\mathcal{A}}(-) := \{(a_1, \dots, a_n) \mid \mathcal{A} \models \varphi[a_1, \dots, a_n]\}$$

be the set of n -tuples *defined by φ in \mathcal{A}* . For $n = 0$ this should be read as⁶

$$(3) \quad \varphi^{\mathcal{A}} := \begin{cases} \text{TRUE} & \text{if } \mathcal{A} \models \varphi \\ \text{FALSE} & \text{if } \mathcal{A} \not\models \varphi. \end{cases}$$

Using this notation we can rewrite (1) as

$$(4) \quad \text{if } \pi : \mathcal{A} \cong \mathcal{B} \text{ then } \pi(\varphi^{\mathcal{A}}(-)) = \varphi^{\mathcal{B}}(-)$$

where for $X \subseteq A^n$ we set $\pi(X) := \{(\pi(a_1), \dots, \pi(a_n)) \mid (a_1, \dots, a_n) \in X\}$.

Later we are going to study various logics that extend first-order logic. In all of these logics only structural properties, that is, properties invariant under isomorphisms, will be expressible. So the analogues of (2) and (4) will be true. In particular, (2) says that only classes of structures closed under isomorphisms can be axiomatizable in these extensions. Only such classes will be of interest. We therefore agree upon the following convention:

Throughout the book all classes K of structures considered will tacitly be assumed to be closed under isomorphisms, i.e.,

$$\mathcal{A} \in K \text{ and } \mathcal{A} \cong \mathcal{B} \text{ imply } \mathcal{B} \in K.$$

(4) shows that properties expressible in logics correspond to so-called global relations:

⁶ TRUE corresponds to the set $\{\emptyset\}$ (consisting of the “empty sequence” \emptyset) and FALSE to the empty set \emptyset .

Definition 1.0.8 Let K be a class of τ -structures. An n -ary *global relation* Γ on K is a mapping assigning to each $\mathcal{A} \in K$ an n -ary relation $\Gamma(\mathcal{A})$ on \mathcal{A} satisfying

$$\Gamma(\mathcal{A})a_1 \dots a_n \quad \text{iff} \quad \Gamma(\mathcal{B})\pi(a_1) \dots \pi(a_n)$$

for every isomorphism $\pi : \mathcal{A} \cong \mathcal{B}$ and every $a_1, \dots, a_n \in A$. If K is the class of all finite τ -structures, then we just speak of an n -ary *global relation*. \square

Examples 1.0.9 (a) Any formula $\varphi(x_1, \dots, x_n)$ of $\text{FO}[\tau]$ defines the global relation $\mathcal{A} \mapsto \varphi^{\mathcal{A}}(-)$.

(b) The “transitive closure relation” TC is the binary global relation on GRAPH with

$$\text{TC}(\mathcal{G}) := \{(a, b) \mid a, b \in G, \text{ there is a path from } a \text{ to } b\}.$$

(c) For $m \geq 0$, Γ_m is a unary global relation on GRAPH, where

$$\Gamma_m(\mathcal{G}) := \{a \mid \|\{b \in G \mid E^{\mathcal{G}}ab\}\| = m\}$$

is the set of elements of \mathcal{G} of degree m . \square

In the definition of a global relation we also allow the case $n = 0$. There are only two 0-ary relations on a structure, TRUE and FALSE. Often one identifies a 0-ary global relation Γ on K with the class

$$\{\mathcal{A} \in K \mid \Gamma(\mathcal{A}) = \text{TRUE}\}.$$

By this identification, the global relation associated with a first-order sentence φ is the class $\text{Mod}(\varphi)$ of finite models of φ (compare (3)).

An important issue in model theory is the study of properties of classes of structures that are axiomatizable in a given logic \mathcal{L} and, in particular, to determine what classes of structures are axiomatizable and what global relations are definable in \mathcal{L} . However, since we only consider “function-free” vocabularies, how can we examine such problems, say, for the class of groups?

Nearly all the methods and results presented in this book can directly be extended to vocabularies containing function symbols. Moreover, by replacing functions by their graphs one can always pass to function-free vocabularies. We sketch how to get rid of the function symbols in a vocabulary τ . The price is the introduction of a new $(n + 1)$ -ary relation symbol F for every n -ary $f \in \tau$. Let the vocabulary τ^r consist of the relation symbols and constants from τ together with the new relation symbols. Thus τ^r is function-free. For a τ -structure \mathcal{A} , let \mathcal{A}^r be the τ^r -structure obtained from \mathcal{A} by replacing every n -ary function $f^{\mathcal{A}}$ by its graph $F^{\mathcal{A}}$,

$$F^{\mathcal{A}} := \{(a_1, \dots, a_n, f(a_1, \dots, a_n)) \mid a_1, \dots, a_n \in A\}.$$

(So, for example, we look at a group (G, \circ^G, e^G) as the $\{R, e\}$ -structure (G, R^G, e^G) , where the ternary relation R is interpreted as the graph of \circ^G ,

i.e., $R^G = \{(a, b, a \circ^G b) \mid a, b \in G\}$ of \circ^G .) The class of τ^r -structures of the form \mathcal{A}^r is the class of models of the conjunction of the formulas

$$\forall x_1 \dots \forall x_n \exists^{=1} y F x_1 \dots x_n y$$

where $f \in \tau$. For every τ -sentence φ there is a τ^r -sentence φ^r and for every τ^r -sentence ψ there is a τ -sentence ψ^{-r} such that for every τ -structure \mathcal{A} we have

$$\begin{array}{lll} \mathcal{A} \models \varphi & \text{iff} & \mathcal{A}^r \models \varphi^r \\ \mathcal{A} \models \psi^{-r} & \text{iff} & \mathcal{A}^r \models \psi. \end{array}$$

(For example, if $\varphi := \exists x \forall y f(g(y)) = x$ then $\varphi^r = \exists x \forall y \exists u (Gyu \wedge Fux)$, and if $\psi := \forall x \exists y (Fxc \wedge \neg Gcy)$ then $\psi^{-r} = \forall x \exists y (f(x) = c \wedge \neg g(c) = y)$. Note that, in general, $\text{qr}(\varphi^r) > \text{qr}(\varphi)$.) Hence, a class K of τ -structures is the class of models of a first-order sentence iff $K^r := \{\mathcal{A}^r \mid \mathcal{A} \in K\}$ is the class of models of a first-order sentence.

E Relational Databases and Query Languages

Suppose that a database contains the names of the main cities in the world and the pairs (a, b) of such cities such that a given airline offers service from a to b without stopover. We can view the database as a first-order structure, namely as a digraph $\mathcal{G} = (G, E^G)$, where G is the set of cities and $E^G ab$ means that there is a flight without stopover from a to b . Now, first-order logic can be considered as a query language. For example, let

$$\varphi(x, y) := Exy \vee \exists z (Exz \wedge Ezy).$$

If φ is thought of as a query to the database, then the response is the set of pairs (a, b) of cities such that b can be reached from a with at most one stop. (We obtain a global relation if we assign to any database (digraph) the response corresponding to the query φ .)

First-order logic provides a rich class of database queries. However, some plausible queries are not first-order expressible. For example, it is impossible to express the query “Can one fly from x to y ” by a first-order formula such that we get the right answer in all possible databases (digraphs). We thus are led to ask for stronger logics (or, query languages).

The last two sections have revealed a close relationship between classes of structures, global relations, and queries. Depending on the type of problem we are studying and the methods that are involved, we shall use one or the other terminology, even though we mostly use the terminology related to classes of structures.

2. The Ehrenfeucht-Fraïssé Method

The Ehrenfeucht-Fraïssé method is among the few tools of model theory that survive when we restrict our attention to finite structures. We present the method in its game-theoretic, its algebraic, and its logical form (due to Ehrenfeucht, Fraïssé, and Hintikka, respectively). Later we shall see that generalizations are also available for some extensions of first-order logic. The detailed presentation for the case of first-order logic will help to understand these extensions, where in each case we only will indicate the changes that are necessary.

We always refer to a fixed vocabulary τ . As already mentioned in the preliminaries, τ contains only relation symbols and constants.

2.1 Elementary Classes

The Ehrenfeucht-Fraïssé method is a useful tool for showing that a given class of structures or a given global relation is (not) definable in first-order logic. We start with some easy remarks concerning the expressive power of first-order logic in the finite.

Proposition 2.1.1 *Every finite structure can be characterized in first-order logic up to isomorphism, i.e., for every finite structure A there is a sentence φ_A of first-order logic such that for all structures B we have*

$$B \models \varphi_A \quad \text{iff} \quad A \cong B.$$

Proof. Suppose $A = \{a_1, \dots, a_n\}$. Set $\bar{a} = a_1 \dots a_n$. Let

$$\Theta_n := \{\psi \mid \psi \text{ has the form } Rx_1 \dots x_k, x = y, \text{ or } c = x, \\ \text{and variables among } v_1, \dots, v_n\}$$

and

$$\varphi_A := \exists v_1 \dots \exists v_n (\bigwedge \{\psi \mid \psi \in \Theta_n, A \models \psi[\bar{a}]\} \wedge \\ \bigwedge \{\neg\psi \mid \psi \in \Theta_n, A \models \neg\psi[\bar{a}]\} \wedge \forall v_{n+1} (v_{n+1} = v_1 \vee \dots \vee v_{n+1} = v_n)).$$

□

Corollary 2.1.2 *Let K be a class of finite structures. Then there is a set Φ of first-order sentences such that*

$$K = \text{Mod}(\Phi),$$

that is, K is the class of finite models of Φ .

Proof. Let K be a class of finite structures. For each n there is only a finite number of pairwise nonisomorphic structures of cardinality n . Let $\{\mathcal{A}_1, \dots, \mathcal{A}_k\}$ be a maximal subset of K of pairwise nonisomorphic structures of cardinality n . Set

$$\psi_n := (\varphi_{=n} \rightarrow (\varphi_{\mathcal{A}_1} \vee \dots \vee \varphi_{\mathcal{A}_k})),^1$$

where $\varphi_{=n}$ is a first-order sentence expressing “there are exactly n elements” (see 1.B). Then $K = \text{Mod}(\{\psi_n \mid n \geq 1\})$. \square

In many situations we want to know more, namely whether a class K of finite structures or, equivalently, a property of finite structures, is axiomatizable by a *single* first-order sentence, i.e., whether K is elementary in the sense of the following definition.

Definition 2.1.3 Let K be a class of finite structures. K is called *axiomatizable in first-order logic* or *elementary*, if there is a sentence φ of first-order logic such that $K = \text{Mod}(\varphi)$.² \square

For structures \mathcal{A} and \mathcal{B} and $m \in \mathbb{N}$ we write $\mathcal{A} \equiv_m \mathcal{B}$ and say that \mathcal{A} and \mathcal{B} are *m -equivalent*, if \mathcal{A} and \mathcal{B} satisfy the same first-order sentences of quantifier rank $\leq m$. The following theorem contains necessary and – as we shall see in 2.2.12 – sufficient conditions for a class K to be elementary. Since, in general, it is used to prove nonaxiomatizability results (see 2.3.5–2.3.9), we formulate it in the corresponding way.

Theorem 2.1.4 *Let K be a class of finite structures. Suppose that for every m there are finite structures \mathcal{A} and \mathcal{B} such that*

$$\mathcal{A} \in K, \mathcal{B} \notin K, \text{ and } \mathcal{A} \equiv_m \mathcal{B}.$$

Then K is not axiomatizable in first-order logic.

Proof. Let φ be any first-order sentence. Set $m := \text{qr}(\varphi)$. By our assumption there are \mathcal{A} and \mathcal{B} such that $\mathcal{A} \in K$, $\mathcal{B} \notin K$, and $\mathcal{A} \equiv_m \mathcal{B}$; hence, $K \neq \text{Mod}(\varphi)$. \square

¹ In case $k = 0$ recall our convention that the empty disjunction is F; one also could set $\psi_n := (\varphi_{=n} \rightarrow \neg \exists x x = x)$ in this case.

² In the literature, instead of *axiomatizable* one often uses the term *finitely axiomatizable*. In (general) model theory a class of *arbitrary* structures K is called *elementary* if, for some φ , K is the class of arbitrary models of φ . And given classes K_0 and K with $K_0 \supseteq K$, it is said that K is *elementary relative to K_0* if, for some φ , K is the class of models of φ in K_0 . In this terminology our notion of *elementary* corresponds to *elementary relative to the class of finite structures*.

2.2 Ehrenfeucht's Theorem

In this section we present a purely gametheoretic characterization of the relation \equiv_m . It will be a useful tool, in particular, for applying Theorem 2.1.4 to concrete classes, at the same time helping us to understand the expressive power of first-order logic. One of the central ingredients of the characterization are partial isomorphisms.

Definition 2.2.1 Assume \mathcal{A} and \mathcal{B} are structures. Let p be a map with $\text{do}(p) \subseteq A$ and $\text{rg}(p) \subseteq B$, where $\text{do}(p)$ and $\text{rg}(p)$ denote the domain and the range of p , respectively. Then p is said to be a *partial isomorphism* from \mathcal{A} to \mathcal{B} if

- p is injective
- for every $c \in \tau$: $c^A \in \text{do}(p)$ and $p(c^A) = c^B$
- for every n -ary $R \in \tau$ and all $a_1, \dots, a_n \in \text{do}(p)$,

$$R^A a_1 \dots a_n \quad \text{iff} \quad R^B p(a_1) \dots p(a_n).$$

We write $\text{Part}(\mathcal{A}, \mathcal{B})$ for the set of partial isomorphisms from \mathcal{A} to \mathcal{B} . \square

In the following we identify a map p with its graph $\{(a, p(a)) \mid a \in \text{do}(p)\}$. Then $p \subseteq q$ means that q is an extension of p .

Remarks 2.2.2 (a) The empty map, $p = \emptyset$, is a partial isomorphism from \mathcal{A} to \mathcal{B} just in case the vocabulary contains no constants.

(b) If $p \neq \emptyset$ is a map with $\text{do}(p) \subseteq A$ and $\text{rg}(p) \subseteq B$, then p is a partial isomorphism from \mathcal{A} to \mathcal{B} iff $\text{do}(p)$ contains c^A for all constants $c \in \tau$ and $p : \text{do}(p)^A \cong \text{rg}(p)^B$ (where $\text{do}(p)^A$ and $\text{rg}(p)^B$ denote the substructures of \mathcal{A} and \mathcal{B} with universes $\text{do}(p)$ and $\text{rg}(p)$, respectively).

(c) For $\bar{a} = a_1 \dots a_s \in A$ and $\bar{b} = b_1 \dots b_s \in B$ the following statements are equivalent:

- (i) The clauses

$$p(a_i) = b_i \text{ for } i = 1, \dots, s$$

and

$$p(c^A) = c^B \text{ for } c \text{ in } \tau$$

define a map, which is a partial isomorphism from \mathcal{A} to \mathcal{B} (henceforth denoted by $\bar{a} \mapsto \bar{b}$, a notation that suppresses the constants).

- (ii) For all quantifier-free $\varphi(v_1, \dots, v_s)$: $\mathcal{A} \models \varphi[\bar{a}]$ iff $\mathcal{B} \models \varphi[\bar{b}]$.
 (iii) For all atomic $\varphi(v_1, \dots, v_s)$: $\mathcal{A} \models \varphi[\bar{a}]$ iff $\mathcal{B} \models \varphi[\bar{b}]$.

Proof. Note that for an arbitrary structure \mathcal{D} and \bar{d} in D ,

$$\begin{aligned} d_i = d_j & \quad \text{iff} \quad \mathcal{D} \models v_i = v_j[\bar{d}] \\ c^D = d_j & \quad \text{iff} \quad \mathcal{D} \models c = v_j[\bar{d}] \\ R^D c^D d_i d_j & \quad \text{iff} \quad \mathcal{D} \models R c v_i v_j[\bar{d}] \end{aligned}$$

($c, R \in \tau, R$ ternary). Using such equivalences, it is easy to show that (i) and (iii) are equivalent. Clearly, (ii) implies (iii), and (ii) follows from (iii), since every quantifier-free formula is a boolean combination of atomic formulas. \square

In general, a partial isomorphism does not preserve the validity of formulas with quantifiers: Let $\tau = \{<\}, \mathcal{A} = (\{0, 1, 2\}, <), \mathcal{B} = (\{0, 1, 2, 3\}, <)$, where in both cases $<$ denotes the natural ordering. Then $p_0 := 02 \mapsto 01$ is a partial isomorphism from \mathcal{A} to \mathcal{B} such that

$$\mathcal{A} \models \exists v_3(v_1 < v_3 \wedge v_3 < v_2) [0, 2]$$

but

$$\mathcal{B} \not\models \exists v_3(v_1 < v_3 \wedge v_3 < v_2) [p_0(0), p_0(2)].$$

Since

$$\mathcal{A} \models (v_1 < v_3 \wedge v_3 < v_2) [0, 2, 1]$$

we see that, for any $p \in \text{Part}(\mathcal{A}, \mathcal{B})$ with $\text{do}(p) = \{0, 2\}$, the validity of

$$\mathcal{B} \models \exists v_3(v_1 < v_3 \wedge v_3 < v_2) [p(0), p(2)]$$

is equivalent to the existence of some $q \in \text{Part}(\mathcal{A}, \mathcal{B})$ which extends p and has 1 in its domain.

This example indicates that the truth of formulas with quantifiers is preserved under partial isomorphisms provided they admit certain extensions. It embodies the basic idea behind the algebraic characterization of \equiv_m we have in mind: The m -equivalence of structures amounts to the existence of partial isomorphisms that can be extended m times. In the gametheoretic terms introduced by the next definition, $\bar{a} \mapsto \bar{b}$ is such a partial isomorphism if the duplicator has a winning strategy for the Ehrenfeucht game $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.

Let \mathcal{A} and \mathcal{B} be τ -structures, $\bar{a} \in \mathcal{A}^s, \bar{b} \in \mathcal{B}^s$, and $m \in \mathbb{N}$. The *Ehrenfeucht game* $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ is played by two players called the *spoiler* and the *duplicator*. Each player has to make m moves in the course of a play. The players take turns. In his i -th move the spoiler first selects a structure, \mathcal{A} or \mathcal{B} , and an element in this structure. If the spoiler chooses e_i in \mathcal{A} then the duplicator in his i -th move must choose an element f_i in \mathcal{B} . If the spoiler chooses f_i in \mathcal{B} then the duplicator must choose an element e_i in \mathcal{A} .

	\mathcal{A}, \bar{a}	\mathcal{B}, \bar{b}
first move	e_1	f_1
second move	e_2	f_2
\vdots	\vdots	\vdots
m -th move	e_m	f_m

As illustrated by this table, at the end elements e_1, \dots, e_m in \mathcal{A} and f_1, \dots, f_m in \mathcal{B} have been chosen. The duplicator *wins* iff $\bar{a}\bar{e} \mapsto \bar{b}\bar{f} \in \text{Part}(\mathcal{A}, \mathcal{B})$ (in case $m = 0$ we just require that $\bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B})$). Otherwise, the spoiler wins.

Equivalently, the spoiler wins if, after some $i \leq m$, $\bar{a}e_1 \dots e_i \mapsto \bar{b}f_1 \dots f_i$ is not a partial isomorphism. We say that a player, the spoiler or the duplicator, has a *winning strategy* in $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, or shortly, that he *wins* $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, if it is possible for him to win each play whatever choices are made by the opponent. We omit a formal definition of the notion of a winning strategy; it is implicit in the algebraic characterization given below (cf. the proof of Theorem 2.3.3). If $s = 0$ (and hence \bar{a} and \bar{b} are empty), we denote the game by $G_m(\mathcal{A}, \mathcal{B})$.

Lemma 2.2.3 (a) *If $\mathcal{A} \cong \mathcal{B}$ then the duplicator wins $G_m(\mathcal{A}, \mathcal{B})$.*

(b) *If the duplicator wins $G_{m+1}(\mathcal{A}, \mathcal{B})$ and $\|A\| \leq m$ then $\mathcal{A} \cong \mathcal{B}$.*

Proof. (a) Suppose $\pi : \mathcal{A} \cong \mathcal{B}$. A winning strategy for the duplicator consists in always choosing the image or preimage under π of the spoiler's selection; that is, if the spoiler chooses $a \in A$ then the duplicator chooses $\pi(a)$; and if the spoiler chooses $b \in B$ then the duplicator answers with $\pi^{-1}(b)$.

(b) Suppose that the duplicator wins $G_{m+1}(\mathcal{A}, \mathcal{B})$, and assume that $A = \{a_1, \dots, a_m\}$. Consider a play where the spoiler, in his first m moves, chooses a_1, \dots, a_m . Let b_1, \dots, b_m be the responses of the duplicator according to his winning strategy. Then $p : \bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B})$ with $\text{do}(p) = A$. It even is an isomorphism from \mathcal{A} onto \mathcal{B} . Otherwise, we have $\text{rg}(p) \neq B$. Then the spoiler, in the last move of the play, chooses some element $b \in B \setminus \text{rg}(p)$. As there is no answer for the duplicator leading to a win, we get a contradiction. \square

The following lemma collects some facts about the Ehrenfeucht game. Their proofs are immediate from the definition.

Lemma 2.2.4 *Let \mathcal{A} and \mathcal{B} be structures, $\bar{a} \in A^s$, $\bar{b} \in B^s$, and $m \geq 0$.*

(a) *The duplicator wins $G_0(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ iff $\bar{a} \mapsto \bar{b}$ is a partial isomorphism.*

(b) *For $m > 0$ the following are equivalent:*

(i) *The duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.*

(ii) *For all $a \in A$ there is $b \in B$ such that the duplicator wins the game $G_{m-1}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$ and for all $b \in B$ there is $a \in A$ such that the duplicator wins $G_{m-1}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$.*

(c) *If the duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ and if $m' < m$, the duplicator wins $G_{m'}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. \square*

Parts (a) and (b) give us a hint how to relate the game to the validity of formulas.

Let \mathcal{A} be given. For $\bar{a} = a_1 \dots a_s \in A$ and $m \geq 0$ we introduce a formula $\varphi_{\bar{a}}^m(v_1, \dots, v_s)$ that describes the gametheoretic properties of \bar{a} in any game $G_m(\mathcal{A}, \bar{a}, \dots)$; more precisely, we want to define $\varphi_{\bar{a}}^m$ in such a way that for any \mathcal{B} and $\bar{b} = b_1 \dots b_s \in B$,

$$\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}] \quad \text{iff} \quad \text{the duplicator wins } G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}).$$

If the structure \mathcal{A} is not clear from the context, we use the notation $\varphi_{\mathcal{A},\bar{a}}^m$ for $\varphi_{\bar{a}}^m$. We also allow $s = 0$, the case of the empty sequence \emptyset of elements in A , and write $\varphi_{\mathcal{A}}^m$ for the sentence $\varphi_{\mathcal{A},\emptyset}^m$.

Definition 2.2.5 Let \bar{v} be v_1, \dots, v_s .

$$\varphi_{\bar{a}}^0(\bar{v}) := \bigwedge \{ \varphi(\bar{v}) \mid \varphi \text{ atomic or negated atomic, } \mathcal{A} \models \varphi[\bar{a}] \}$$

and for $m > 0$,

$$\varphi_{\bar{a}}^m(\bar{v}) := \bigwedge_{a \in A} \exists v_{s+1} \varphi_{\bar{a}a}^{m-1}(\bar{v}, v_{s+1}) \wedge \forall v_{s+1} \bigvee_{a \in A} \varphi_{\bar{a}a}^{m-1}(\bar{v}, v_{s+1}). \quad \square$$

$\varphi_{\bar{a}}^0$ describes the isomorphism type of the substructure generated by \bar{a} in \mathcal{A} ; and for $m > 0$ the formula $\varphi_{\bar{a}}^m$ tells us to which isomorphism types the tuple \bar{a} can be extended in m steps adding one element in each step. $\varphi_{\bar{a}}^m$ is called the *m-isomorphism type* (or *m-Hintikka formula*) of \bar{a} in \mathcal{A} .

Since $\{ \varphi(v_1, \dots, v_s) \mid \varphi \text{ atomic or negated atomic} \}$ is finite, a simple induction on m shows:

Lemma 2.2.6 For $s, m \geq 0$, the set $\{ \varphi_{\mathcal{A},\bar{a}}^m \mid \mathcal{A} \text{ a structure and } \bar{a} \in A^s \}$ is finite. \square

In particular, the conjunctions and disjunctions in the above definition are finite.

Lemma 2.2.7 (a) $\text{qr}(\varphi_{\bar{a}}^m) = m$.

(b) $\mathcal{A} \models \varphi_{\bar{a}}^m[\bar{a}]$.

(c) For any \mathcal{B} and \bar{b} in B ,

$$\mathcal{B} \models \varphi_{\bar{a}}^0[\bar{b}] \quad \text{iff} \quad \bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B}).$$

Proof. The proofs of (a) and (b) are straightforward. (c) holds by part (c) of 2.2.2. \square

Theorem 2.2.8 (Ehrenfeucht's Theorem) Given \mathcal{A} and \mathcal{B} , $\bar{a} \in A^s$ and $\bar{b} \in B^s$, and $m \geq 0$, the following are equivalent:

(i) The duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.

(ii) $\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}]$.

(iii) \bar{a} and \bar{b} satisfy the same formulas of quantifier rank $\leq m$, that is, if $\varphi(x_1, \dots, x_s)$ is of quantifier rank $\leq m$, then

$$(*) \quad \mathcal{A} \models \varphi[\bar{a}] \quad \text{iff} \quad \mathcal{B} \models \varphi[\bar{b}].$$

Proof. (iii) implies (ii) since $\text{qr}(\varphi_{\bar{a}}^m) = m$ and $\mathcal{A} \models \varphi_{\bar{a}}^m[\bar{a}]$. We prove the equivalence of (i) and (ii) by induction on m . For $m = 0$

$$\begin{aligned} \text{the duplicator wins } G_0(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}) & \quad \text{iff} \quad \bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B}) \quad (\text{cf. 2.2.4(a)}) \\ & \quad \text{iff} \quad \mathcal{B} \models \varphi_{\bar{a}}^0[\bar{b}] \quad (\text{by 2.2.7(c)}). \end{aligned}$$

For $m > 0$,

- the duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$
- iff for all $a \in A$ there is $b \in B$ such that the duplicator wins $G_{m-1}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$, and for all $b \in B$ there is $a \in A$ such that the duplicator wins $G_{m-1}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$ (cf. 2.2.4(b))
- iff for all $a \in A$ there is $b \in B$ with $\mathcal{B} \models \varphi_{\bar{a}a}^{m-1}[\bar{b}b]$, and for all $b \in B$ there is $a \in A$ with $\mathcal{B} \models \varphi_{\bar{a}a}^{m-1}[\bar{b}b]$ (ind. hyp.)
- iff $\mathcal{B} \models \bigwedge_{a \in A} \exists v_{s+1} \varphi_{\bar{a}a}^{m-1}(\bar{v}, v_{s+1}) \wedge \forall v_{s+1} \bigvee_{a \in A} \varphi_{\bar{a}a}^{m-1}(\bar{v}, v_{s+1})[\bar{b}]$
- iff $\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}]$.

(i) \Rightarrow (iii): The proof proceeds by induction on m . The case $m = 0$ is handled as above. Let $m > 0$ and suppose that the duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. Clearly, the set of formulas $\varphi(x_1, \dots, x_s)$ satisfying $(*)$ contains the atomic formulas and is closed under \neg and \vee . Suppose that $\varphi(\bar{x}) = \exists y\psi$ and $\text{qr}(\varphi) \leq m$. Since $y \notin \text{free}(\varphi)$, we can assume that y is distinct from the variables in \bar{x} . Hence, $\psi = \psi(\bar{x}, y)$. Assume, for instance, $\mathcal{A} \models \varphi[\bar{a}]$. Then there is $a \in A$ such that $\mathcal{A} \models \psi[\bar{a}, a]$. As, by (i), the duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, there is $b \in B$ such that the duplicator wins $G_{m-1}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$. Since $\text{qr}(\psi) \leq m-1$, the induction hypothesis yields $\mathcal{B} \models \psi[\bar{b}, b]$, hence $\mathcal{B} \models \varphi[\bar{b}]$. \square

Corollary 2.2.9 *For structures \mathcal{A}, \mathcal{B} and $m \geq 0$ the following are equivalent:*

- (i) *The duplicator wins $G_m(\mathcal{A}, \mathcal{B})$.*
- (ii) $\mathcal{B} \models \varphi_{\bar{a}}^m$.
- (iii) $\mathcal{A} \equiv_m \mathcal{B}$. \square

By 2.2.3(b) we get

Corollary 2.2.10 *Let \mathcal{A} be a structure with $\|\mathcal{A}\| \leq m$. Then for all \mathcal{B} ,*

$$\mathcal{B} \models \varphi_{\bar{a}}^{m+1} \quad \text{iff} \quad \mathcal{A} \cong \mathcal{B}. \quad \square$$

The next result shows that the formulas $\varphi_{\bar{a}}^m$ give a clear picture of the expressive power of first-order logic.

Theorem 2.2.11 *Let $\varphi(v_1, \dots, v_s)$ be a formula of quantifier rank $\leq m$. Then*

$$\models \varphi \leftrightarrow \bigvee \{ \varphi_{\bar{a}}^m \mid \mathcal{A} \text{ a structure, } \bar{a} \in \mathcal{A}, \text{ and } \mathcal{A} \models \varphi[\bar{a}] \}.$$
³

Proof. Suppose first that $\mathcal{B} \models \varphi[\bar{b}]$. Then the formula $\varphi_{\bar{b}}^m$ is a member of the disjunction on the right side of the equivalence, which therefore is satisfied by \bar{b} . Conversely, suppose $\mathcal{B} \models \bigvee \{ \varphi_{\bar{a}}^m \mid \mathcal{A} \models \varphi[\bar{a}] \}[\bar{b}]$. Then, for some \mathcal{A} and \bar{a} such that $\mathcal{A} \models \varphi[\bar{a}]$, we have $\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}]$. By 2.2.8, \bar{a} and \bar{b} satisfy the same formulas of quantifier rank $\leq m$ and therefore, $\mathcal{B} \models \varphi[\bar{b}]$. \square

³ By 2.2.6 the disjunction is taken over a finite set.

The following result contains the desired characterization of classes axiomatizable in first-order logic. As 2.1.4, we formulate it in a “negative” form, since the implication (ii) \Rightarrow (i) turns out to be one of the main tools to obtain nonaxiomatizability results for first-order logic.

Theorem 2.2.12 *For a class K of finite structures the following are equivalent:*

- (i) K is not axiomatizable in first-order logic.
- (ii) For each m there are finite structures \mathcal{A} and \mathcal{B} such that

$$\mathcal{A} \in K, \mathcal{B} \notin K \text{ and } \mathcal{A} \equiv_m \mathcal{B}.$$

Proof. (ii) \Rightarrow (i) was proven in 2.1.4. For the converse, suppose that (ii) does not hold, i.e., that for some m and all finite \mathcal{A} and \mathcal{B} ,

$$\mathcal{A} \in K \text{ and } \mathcal{A} \equiv_m \mathcal{B} \text{ imply } \mathcal{B} \in K.$$

Then $K = \text{Mod}(\bigvee \{\varphi_{\mathcal{A}}^m \mid \mathcal{A} \in K\})$, and thus K is axiomatizable. □

2.3 Examples and Fraïssé’s Theorem

To apply Ehrenfeucht’s characterization to concrete examples, it is more convenient to use an algebraic version due to Fraïssé, even though it lacks the intuitive appeal of the gametheoretical approach.

Given structures \mathcal{A}, \mathcal{B} and $m \in \mathbb{N}$, let $W_m(\mathcal{A}, \mathcal{B}) :=$

$$\{\bar{a} \mapsto \bar{b} \mid s \geq 0, \bar{a} \in A^s, \bar{b} \in B^s, \text{ the duplicator wins } G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})\}$$

be the set of winning positions for the duplicator. The sequence of the $W_m(\mathcal{A}, \mathcal{B})$ has the back and forth properties as introduced in the following definition.

Definition 2.3.1 Structures \mathcal{A} and \mathcal{B} are said to be *m -isomorphic*, written $\mathcal{A} \cong_m \mathcal{B}$, if there is a sequence $(I_j)_{j \leq m}$ with the following properties:

- (a) Every I_j is a nonempty set of partial isomorphisms from \mathcal{A} to \mathcal{B} .
- (b) (*Forth property*) For every $j < m$, $p \in I_{j+1}$, and $a \in A$ there is $q \in I_j$ such that $q \supseteq p$ and $a \in \text{do}(q)$.
- (c) (*Back property*) For every $j < m$, $p \in I_{j+1}$, and $b \in B$ there is $q \in I_j$ such that $q \supseteq p$ and $b \in \text{rg}(q)$.

If $(I_j)_{j \leq m}$ has the properties (a), (b), and (c), we write $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$ and say that \mathcal{A} and \mathcal{B} are *m -isomorphic via $(I_j)_{j \leq m}$* . □

Exercise 2.3.2 Suppose $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$. Then $(\tilde{I}_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$ with $\tilde{I}_j := \{q \in \text{Part}(\mathcal{A}, \mathcal{B}) \mid q \subseteq p \text{ for some } p \in I_j\}$. In particular, $\emptyset \mapsto \emptyset \in I_j$ for all $j \leq m$. Moreover, $W_j(\mathcal{A}, \mathcal{B}) = \widetilde{W}_j(\mathcal{A}, \mathcal{B})$.

Using the results of the preceding section we obtain:

Theorem 2.3.3 *For structures \mathcal{A} and \mathcal{B} , $\bar{a} \in A^s$, $\bar{b} \in B^s$, and $m \geq 0$ the following are equivalent:*

- (i) *The duplicator wins $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.*
- (ii) *$\bar{a} \mapsto \bar{b} \in W_m(\mathcal{A}, \mathcal{B})$ and $(W_j(\mathcal{A}, \mathcal{B}))_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$.*
- (iii) *There is $(I_j)_{j \leq m}$ with $\bar{a} \mapsto \bar{b} \in I_m$ such that $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$.*
- (iv) *$\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}]$.*
- (v) *\bar{a} satisfies in \mathcal{A} the same formulas of quantifier rank $\leq m$ as \bar{b} in \mathcal{B} .*

Proof. By the definition of $W_m(\mathcal{A}, \mathcal{B})$ and 2.2.4, (i) implies (ii). Obviously, (ii) implies (iii). Therefore it suffices to show the implication (iii) \Rightarrow (i), the remaining equivalences being clear from 2.2.8. For (iii) \Rightarrow (i) suppose that $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$ and $\bar{a} \mapsto \bar{b} \in I_m$. We describe a winning strategy in $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ for the duplicator: In his i -th move he should choose the element e_i (or f_i , respectively) such that for $p_i : \bar{a}e_1 \dots e_i \mapsto \bar{b}f_1 \dots f_i$ it is true that $p_i \subseteq q$ for some $q \in I_{m-i}$; this is always possible because of the back and forth properties of $(I_j)_{j \leq m}$. Looking at $i := m$ we see that the duplicator wins. \square

For $s = 0$ in view of Exercise 2.3.2 the preceding theorem yields the following extension of 2.2.9.

Corollary 2.3.4 *For structures \mathcal{A}, \mathcal{B} and $m \geq 0$ the following are equivalent:*

- (i) *The duplicator wins $G_m(\mathcal{A}, \mathcal{B})$.*
- (ii) *$(W_j(\mathcal{A}, \mathcal{B}))_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$.*
- (iii) *$\mathcal{A} \cong_m \mathcal{B}$.*
- (iv) *$\mathcal{B} \models \varphi_{\bar{a}}^m$.*
- (v) *$\mathcal{A} \equiv_m \mathcal{B}$.*

\square

The equivalence of (iii) and (v) is known as *Fraïssé's Theorem*. The proof of the preceding theorem, and especially that of the equivalence of (i) and (iii) shows, that Ehrenfeucht's Theorem and Fraïssé's Theorem are different formulations of the same fact. In particular, the proof exhibits the close relationship between sequences $(I_j)_{j \leq m}$ and winning strategies for the duplicator in $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. Therefore, one often speaks of the *Ehrenfeucht-Fraïssé game* or the *Ehrenfeucht-Fraïssé method*.

Example 2.3.5 Let τ be the empty vocabulary and \mathcal{A} and \mathcal{B} be τ -structures (i.e., nonempty sets). Suppose $\|\mathcal{A}\| \geq m$ and $\|\mathcal{B}\| \geq m$. Then $\mathcal{A} \cong_m \mathcal{B}$. In fact, $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$ with $I_j := \{p \in \text{Part}(\mathcal{A}, \mathcal{B}) \mid \|\text{do}(p)\| \leq m - j\}$.

As a consequence the class $\text{EVEN}[\tau]$ of finite τ -structures of even cardinality is not axiomatizable in first-order logic. In fact, for each $m > 0$, let \mathcal{A}_m be a structure of cardinality m . Then, $\mathcal{A}_m \in \text{EVEN}[\tau]$ iff $\mathcal{A}_{m+1} \notin \text{EVEN}[\tau]$, but $\mathcal{A}_m \cong_m \mathcal{A}_{m+1}$. Now apply 2.2.12. The reader is encouraged to show for arbitrary τ that $\text{EVEN}[\tau]$ is not axiomatizable. \square

Example 2.3.6 Let $\tau = \{<, \min, \max\}$ be a vocabulary for finite orderings as introduced in 1.A2 and $m \geq 1$. Suppose that \mathcal{A} and \mathcal{B} are finite orderings, $\|A\| > 2^m$ and $\|B\| > 2^m$. Then $\mathcal{A} \cong_m \mathcal{B}$. Hence, the class of finite orderings of even cardinality is not axiomatizable in first-order logic. Clearly, the last statement remains true, if we consider orderings as $\{<, S, \min, \max\}$ -structures.

For a *proof*, given any ordering C , we define its distance function d by

$$d(a, a') := \|\{b \in C \mid (a < b \leq a') \text{ or } (a' < b \leq a)\}\|.$$

And, for $j \geq 0$, we introduce the “truncated” j -distance function d_j on $C \times C$ by

$$d_j(a, a') := \begin{cases} d(a, a') & \text{if } d(a, a') < 2^j \\ \infty & \text{else.} \end{cases}$$

Now, suppose that \mathcal{A} and \mathcal{B} are finite orderings with $\|A\|, \|B\| > 2^m$. For $j \leq m$ set

$$I_j := \{p \in \text{Part}(\mathcal{A}, \mathcal{B}) \mid d_j(a, a') = d_j(p(a), p(a')) \text{ for } a, a' \in \text{do}(p)\}.$$

Then $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$: By assumption on the cardinalities of \mathcal{A} and \mathcal{B} we have $\{(\min^A, \min^B), (\max^A, \max^B)\} \in I_j$ for every $j \leq m$. To give a proof of the forth property of $(I_j)_{j \leq m}$ (the back property can be proven analogously), suppose $j < m$, $p \in I_{j+1}$, and $a \in A$. We distinguish two cases, depending on whether or not the following condition



(*) there is an $a' \in \text{do}(p)$ such that $d_j(a, a') < 2^j$

is satisfied. If (*) holds then there is exactly one $b \in B$ for which $p \cup \{(a, b)\}$ is a partial isomorphism preserving d_j -distances. Now assume that (*) does not hold and let $\text{do}(p) = \{a_1, \dots, a_r\}$ with $a_1 < \dots < a_r$. We restrict ourselves to the case $a_i < a < a_{i+1}$ for some i . Then, $d_j(a_i, a) = \infty$ and $d_j(a, a_{i+1}) = \infty$; hence, $d_{j+1}(a_i, a_{i+1}) = \infty$ and therefore, $d_{j+1}(p(a_i), p(a_{i+1})) = \infty$. Thus there is a b such that $p(a_i) < b < p(a_{i+1})$, $d_j(p(a_i), b) = \infty$, and $d_j(b, p(a_{i+1})) = \infty$. One easily verifies that $q := p \cup \{(a, b)\}$ is a partial isomorphism in I_j . \square

Example 2.3.7 Let $\tau = \{<, \min, \max\}$ be as in the preceding example and $\sigma = \tau \cup \{E\}$ with a binary relation symbol E . For $n \geq 3$ let \mathcal{A}_n be the ordered τ -structure with $A_n = \{0, \dots, n\}$, $\min^{A_n} = 0$, $\max^{A_n} = n$, where $<^{A_n}$ is the natural ordering on A_n , and

$$E^{A_n} = \{(i, j) \mid |i - j| = 2\} \cup \{(0, n), (n, 0), (1, n-1), (n-1, 1)\}.$$

(A_n, E^{A_n}) is a graph that is connected iff n is odd. Now, let $m \geq 2$ and $l, k \geq 2^m$.



Let I_j be the set of partial isomorphisms from $\mathcal{A}_l|_\tau$ to $\mathcal{A}_k|_\tau$ as introduced in the preceding example. For $j \geq 2$ any $p \in I_j$ preserves E , too, that is, $I_j \subseteq \text{Part}(\mathcal{A}_l, \mathcal{A}_k)$. Hence, $(I_{j+2})_{j \leq m-2} : \mathcal{A}_l \cong_{m-2} \mathcal{A}_k$, and, by 2.2.12, we obtain:

- The class of finite connected ordered graphs is not first-order axiomatizable. □

The following example is related to the preceding one, but does not involve orderings.

Example 2.3.8 For $l \geq 1$, let \mathcal{G}_l be the graph given by a cycle of length $l + 1$. To be precise, set

$$\mathcal{G}_l := \{0, \dots, l\}, \quad E^{\mathcal{G}_l} := \{(i, i + 1) \mid i < l\} \cup \{(i + 1, i) \mid i < l\} \cup \{(0, l), (l, 0)\}.$$

Thus, for $l, k \in \mathbb{N}$, the disjoint union $\mathcal{G}_l \dot{\cup} \mathcal{G}_k$ consists of a cycle of length $l + 1$ and of a cycle of length $k + 1$ (for \mathcal{A}_l as defined in the preceding example we have $\mathcal{A}_l|_{\{E\}} \cong \mathcal{G}_l$ for l odd, and $\mathcal{A}_l|_{\{E\}} \cong \mathcal{G}_{\frac{l}{2}-1} \dot{\cup} \mathcal{G}_{\frac{l}{2}}$ for l even). We show:

$$\text{If } l, k \geq 2^m \text{ then } \mathcal{G}_l \cong_m \mathcal{G}_k \text{ and } \mathcal{G}_l \cong_m \mathcal{G}_l \dot{\cup} \mathcal{G}_l.$$

In fact, for $j \in \mathbb{N}$, define the distance function d_j on a graph \mathcal{G} by

$$d_j(a, a') := \begin{cases} d(a, b) & \text{if } d(a, b) < 2^{j+1} \\ \infty & \text{else} \end{cases}$$

(where d denotes the distance function on \mathcal{G} introduced in 1.A1). To show, say, that \mathcal{G}_l and $\mathcal{G}_l \dot{\cup} \mathcal{G}_l$ are m -isomorphic, one verifies $(I_j)_{j \leq m} : \mathcal{G}_l \cong_m \mathcal{G}_l \dot{\cup} \mathcal{G}_l$ where I_j is the set of $p \in \text{Part}(\mathcal{G}_l, \mathcal{G}_l \dot{\cup} \mathcal{G}_l)$ with

$$\|\text{do}(p)\| \leq m - j \quad \text{and} \quad d_j(a, b) = d_j(p(a), p(b)) \text{ for } a, b \in \text{do}(p)$$

(the proof is similar to that of Example 2.3.6). We note two consequences.

- The class CONN of connected finite graphs is not axiomatizable in first-order logic.

In fact, by 2.2.12, CONN is not axiomatizable, since for each m we have

$$\mathcal{G}_{2^m} \in \text{CONN}, \quad \mathcal{G}_{2^m} \dot{\cup} \mathcal{G}_{2^m} \notin \text{CONN}, \quad \mathcal{G}_{2^m} \equiv_m \mathcal{G}_{2^m} \dot{\cup} \mathcal{G}_{2^m}.$$

– The global relation TC (cf. 1.0.9), the relation of transitive closure on the class GRAPH of finite graphs, is not first-order definable.

In fact, suppose $\psi(x, y)$ is a first-order formula defining TC on GRAPH. Then CONN would be the class of finite models of $\forall x \forall y (\neg x = y \rightarrow \psi(x, y))$ (and the graph axioms). \square

Exercise 2.3.9 Set $\tau = \{E\}$. For $l \geq 1$ let \mathcal{B}_l and \mathcal{D}_l be the τ -structures given by

$$\begin{aligned} B_l &:= \{0, \dots, l\}, & E^{B_l} &:= \{(i, i + 1) \mid i < l\}, \\ D_l &:= \{0, \dots, l\}, & E^{D_l} &:= \{(i, i + 1) \mid i < l\} \cup \{(l, 0)\}. \end{aligned}$$

Given $m \geq 0$, show that $\mathcal{B}_l \cong_m \mathcal{B}_l \dot{\cup} \mathcal{D}_l$ for sufficiently large l . Conclude that the class of finite acyclic digraphs (cf. 1.A1) is not axiomatizable in first-order logic. \square

Proposition 2.3.10 *The product, the disjoint union, and the ordered sum (cf. 1.A3) preserve \equiv_m , i.e.,*

- (a) *If $\mathcal{A}_1 \equiv_m \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m \mathcal{B}_2$ then $\mathcal{A}_1 \times \mathcal{A}_2 \equiv_m \mathcal{B}_1 \times \mathcal{B}_2$.*
- (b) *If $\mathcal{A}_1 \equiv_m \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m \mathcal{B}_2$ then $\mathcal{A}_1 \dot{\cup} \mathcal{A}_2 \equiv_m \mathcal{B}_1 \dot{\cup} \mathcal{B}_2$.*
- (c) *If $\mathcal{A}_1 \equiv_m \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m \mathcal{B}_2$ then $\mathcal{A}_1 \triangleleft \mathcal{A}_2 \equiv_m \mathcal{B}_1 \triangleleft \mathcal{B}_2$.*

Proof. Suppose $\mathcal{A}_1 \equiv_m \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m \mathcal{B}_2$. By Ehrenfeucht's Theorem there are winning strategies for the duplicator in the games $G_m(\mathcal{A}_1, \mathcal{B}_1)$ and $G_m(\mathcal{A}_2, \mathcal{B}_2)$. We refer to these strategies as S_1 and S_2 .

(a) The following gives a winning strategy for the duplicator in the game $G_m(\mathcal{A}_1 \times \mathcal{A}_2, \mathcal{B}_1 \times \mathcal{B}_2)$: We simultaneously play games in $G_m(\mathcal{A}_1, \mathcal{B}_1)$ and $G_m(\mathcal{A}_2, \mathcal{B}_2)$. Suppose that in his i -th move the spoiler chooses, say, $(a_1, a_2) \in \mathcal{A}_1 \times \mathcal{A}_2$. Let $b_1 \in \mathcal{B}_1$ and $b_2 \in \mathcal{B}_2$ be answers to a_1 and a_2 according to S_1 and S_2 , respectively. Then the duplicator chooses (b_1, b_2) .

(b),(c) The proofs for parts (b) and (c) proceed in the same way. So let $*$ $\in \{\dot{\cup}, \triangleleft\}$. The following represents a winning strategy for the duplicator in $G_m(\mathcal{A}_1 * \mathcal{A}_2, \mathcal{B}_1 * \mathcal{B}_2)$ (when describing it we use moves of plays in $G_m(\mathcal{A}_1, \mathcal{B}_1)$ and $G_m(\mathcal{A}_2, \mathcal{B}_2)$). Suppose that in his i -th move the spoiler selects, say, $a \in \mathcal{A}_1 * \mathcal{A}_2$. Then the duplicator gets his answer by applying S_1 if $a \in \mathcal{A}_1$, and S_2 if $a \in \mathcal{A}_2$. \square

The last proof yields more:

Corollary 2.3.11 (a) *If $(\mathcal{A}_1, \bar{a}_1) \equiv_m (\mathcal{B}_1, \bar{b}_1)$ and $(\mathcal{A}_2, \bar{a}_2) \equiv_m (\mathcal{B}_2, \bar{b}_2)$ then $(\mathcal{A}_1 \dot{\cup} \mathcal{A}_2, \bar{a}_1, \bar{a}_2) \equiv_m (\mathcal{B}_1 \dot{\cup} \mathcal{B}_2, \bar{b}_1, \bar{b}_2)$.*
 (b) *If $(\mathcal{A}_1, \bar{a}_1) \equiv_m (\mathcal{B}_1, \bar{b}_1)$ and $(\mathcal{A}_2, \bar{a}_2) \equiv_m (\mathcal{B}_2, \bar{b}_2)$ then $(\mathcal{A}_1 \triangleleft \mathcal{A}_2, \bar{a}_1, \bar{a}_2) \equiv_m (\mathcal{B}_1 \triangleleft \mathcal{B}_2, \bar{b}_1, \bar{b}_2)$.⁴ \square*

⁴ In this corollary, \bar{a}_1 and \bar{a}_2 can be sequences of different length.

Exercise 2.3.12 Suppose that $\tau = \{R_1, \dots, R_k\}$ with unary relation symbols R_i . For $\alpha : \{1, \dots, k\} \rightarrow \{0, 1\}$ and a τ -structure \mathcal{A} denote by A_α the subset $X_1 \cap \dots \cap X_k$, where $X_i := R_i^{\mathcal{A}}$ if $\alpha(i) = 1$, and $X_i := A \setminus R_i^{\mathcal{A}}$ otherwise. Show for any τ -structures \mathcal{A}, \mathcal{B} and $m \geq 1$ that $\mathcal{A} \equiv_m \mathcal{B}$ iff $\min\{\|A_\alpha\|, m\} = \min\{\|B_\alpha\|, m\}$ holds for all α . Conclude that every sentence φ from $\text{FO}[\tau]$ is equivalent to a boolean combination of sentences of the form $\exists^{=l} x R^\alpha x$, where $R^\alpha x := \varphi_1 \wedge \dots \wedge \varphi_k$ with $\varphi_i = R_i x$ if $\alpha(i) = 1$, and $\varphi_i = \neg R_i x$ if $\alpha(i) = 0$. Hint: First, consider sentences φ of the form $\varphi_{\mathcal{A}}^m$. Then apply 2.2.11. \square

Exercise 2.3.13 Let $<$ be a binary symbol of the relational vocabulary τ . For $m \geq 1$ and $l, k > 2^m$ show that $<^l \mathcal{A} \cong_m <^k \mathcal{A}$ holds for ordered \mathcal{A} . Here $<^l \mathcal{A}$ denotes the ordered sum of l copies of \mathcal{A} . (Hint: Use the idea of the proof of 2.3.6). \square

Exercise 2.3.14 Show for $\bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B})$ and $m \geq 0$ that

$$\mathcal{B} \models \varphi_{\bar{a}}^m[\bar{b}] \quad \text{iff} \quad (\mathcal{A}, \bar{a}) \cong_m (\mathcal{B}, \bar{b}). \quad \square$$

Exercise 2.3.15 Show that $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$ implies $I_j \subseteq W_j(\mathcal{A}, \mathcal{B})$ for $j \leq m$. \square

Remark 2.3.16 Let $\mathcal{Z} := (\mathbb{Z}, <)$ and $\mathcal{Q} := (\mathbb{Q}, <)$ be the integers and the rationals with their orderings, respectively. For

$$\varphi := \exists x \exists y (x < y \wedge \forall z \neg (x < z \wedge z < y))$$

we have

$$(*) \quad \mathcal{Z} \models \varphi \quad \text{and} \quad \mathcal{Q} \not\models \varphi.$$

Hence, $\mathcal{Z} \not\equiv_3 \mathcal{Q}$ and therefore, $\mathcal{Z} \not\cong_3 \mathcal{Q}$. The spoiler can “transform” the information $(*)$ into a winning strategy for the game $G_3(\mathcal{Z}, \mathcal{Q})$ as given by the table

<u>5</u>	<u>a</u>
6	b
?	<u>$\frac{a+b}{2}$</u>

We have underlined the selections of the spoiler. Note that no third move of the duplicator will lead to a partial isomorphism (since for $a < b$ we have $a < \frac{a+b}{2} < b$ and there is no integer between 5 and 6). In this strategy of the spoiler his selections in \mathcal{Z} correspond to the existential quantifiers in φ and his selections in \mathcal{Q} to the universal quantifiers. This connection can be made precise and is implicit in the proof of 2.2.8. It is the reason why moves in $G_m(\mathcal{A}, \mathcal{B})$ in which the spoiler chooses an element of \mathcal{A} (of \mathcal{B}) are sometimes called \exists -moves (\forall -moves). \square

As in the present section, the formulas $\varphi_{\mathcal{A}, \bar{a}}^m$, the m -isomorphism type of \bar{a} in \mathcal{A} , will also play a crucial role in subsequent considerations. In our opinion their methodological importance stems from the following two facts:

- (1) They have a clear algebraic meaning.
- (2) Every first-order formula is equivalent to a disjunction of such formulas (see 2.2.11).

Classical model theory has been characterized by the equation

$$\text{model theory} = \text{universal algebra} + \text{logic}$$

(see [20]). By (1) and (2) above, it is clear that the formulas $\varphi_{\mathcal{A}, \bar{a}}^j$ provide a bridge between structures and first-order formulas, that is, between our main notions from (universal) algebra and from (first-order) logic, respectively. Therefore the experience that they are a valuable tool in model theory might not come as a surprise.

There is a more algebraic, sort of logic-free, way to define m -isomorphism types, say for $\bar{a} = a_1 \dots a_s$ in \mathcal{A} by

$$\text{IT}^0(\mathcal{A}, \bar{a}) := \{ \varphi \mid \mathcal{A} \models \varphi[\bar{a}], \varphi(v_1, \dots, v_s) \text{ atomic} \}$$

and

$$\text{IT}^{m+1}(\mathcal{A}, \bar{a}) := \{ \text{IT}^m(\mathcal{A}, \bar{a}a) \mid a \in A \}.$$

One easily verifies that for any \mathcal{B} and $\bar{b} \in \mathcal{B}$,

$$\text{IT}^m(\mathcal{A}, \bar{a}) = \text{IT}^m(\mathcal{B}, \bar{b}) \quad \text{iff} \quad \varphi_{\mathcal{A}, \bar{a}}^m = \varphi_{\mathcal{B}, \bar{b}}^m.$$

2.4 Hanf's Theorem

All vocabularies in this and the next section will be relational unless stated otherwise. For a nonempty subset M of a structure \mathcal{A} we denote by \mathcal{M} the substructure of \mathcal{A} with universe M .

Given a structure \mathcal{A} , we define the binary relation E^A on A by

$$E^A := \{ (a, b) \mid a \neq b, \text{ and there are } R \text{ in } \tau \text{ and } \bar{c} \in A \text{ such that } R^A \bar{c} \text{ and } a \text{ and } b \text{ are components of the tuple } \bar{c} \}.$$

The structure $\mathcal{G}(\mathcal{A}) := (A, E^A)$ is called the *Gaifman graph* of \mathcal{A} . Obviously, if \mathcal{A} itself is a graph then $\mathcal{G}(\mathcal{A}) = \mathcal{A}$. For a in A and $r \in \mathbb{N}$ we denote by $S(r, a)$ (or $S^A(r, a)$) the r -ball of a ,

$$S(r, a) := \{ b \in A \mid d(a, b) \leq r \}.$$
⁵

$S(r, a)$ (or $S^A(r, a)$) stands for the substructure of \mathcal{A} with universe $S(r, a)$. Note that for $b, c \in S(r, a)$ we have $d(b, c) \leq 2r$. For $\bar{a} = a_1 \dots a_s$ we set $S(r, \bar{a}) := S(r, a_1) \cup \dots \cup S(r, a_s)$.

⁵ d denotes the distance function of $\mathcal{G}(\mathcal{A})$ as defined in 1.A1.

We define the r -ball type of a point a in \mathcal{A} to be the isomorphism type of $(\mathcal{S}(r, a), a)$, i.e., points a in \mathcal{A} and b in \mathcal{B} have the same r -ball type iff $(\mathcal{S}^{\mathcal{A}}(r, a), a) \cong (\mathcal{S}^{\mathcal{B}}(r, b), b)$.

In 2.3.8 we showed for certain graphs that they are m -isomorphic and hence, m -equivalent, using a sequence $(I_j)_{j \leq m}$, where – in the terminology just introduced – for each $p \in I_j$ and $a \in \text{do}(p)$ there was an isomorphism of $\mathcal{S}(2^j - 1, a)$ onto $\mathcal{S}(2^j - 1, p(a))$ compatible with p . We use generalizations of this idea to show two further theorems on the expressive power of first-order logic. The first one (Hanf's Theorem) is obtained by applying the idea just mentioned to graphs of structures. In the second one (Gaifman's Theorem) the requirement of isomorphism of corresponding balls is weakened to l -equivalence for a suitable l .

Theorem 2.4.1 (Hanf's Theorem) *Let \mathcal{A} and \mathcal{B} be τ -structures and let $m \in \mathbb{N}$. Suppose that for some $e \in \mathbb{N}$ the 3^m -balls in \mathcal{A} and \mathcal{B} have less than e elements⁶ and that for each 3^m -ball type ι , (i) or (ii) holds where*

- (i) \mathcal{A} and \mathcal{B} have the same number of elements of 3^m -ball type ι ,
- (ii) both \mathcal{A} and \mathcal{B} have more than $m \cdot e$ elements of 3^m -ball type ι .

Then $\mathcal{A} \equiv_m \mathcal{B}$.

Proof. Since for $n < l$ the l -ball type of an element determines its n -ball type, we see that for $n < 3^m$ and every n -ball type ι , \mathcal{A} and \mathcal{B} have the same number of elements of n -ball type ι or both, \mathcal{A} and \mathcal{B} , have more than $m \cdot e$ elements of n -ball type ι .

We show that $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$, where I_j is the set

$$\{\bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B}) \mid (\mathcal{S}(3^j, \bar{a}), \bar{a}) \cong (\mathcal{S}(3^j, \bar{b}), \bar{b}) \text{ and } \text{length}(\bar{a}) \leq m - j\},$$

and where, for $\text{length}(\bar{a}) = 0$, we set $(\mathcal{S}(3^j, \bar{a}), \bar{a}) = \emptyset$ and agree that $\emptyset \cong \emptyset$. Therefore, we have $\emptyset \mapsto \emptyset \in I_m$. Concerning the back and forth properties it is enough, by symmetry, to prove the forth property. Thus suppose that $0 \leq j < m$, $a \in A$ and $\bar{a} \mapsto \bar{b} \in I_{j+1}$, say,

$$(*) \quad \pi : (\mathcal{S}(3^{j+1}, \bar{a}), \bar{a}) \cong (\mathcal{S}(3^{j+1}, \bar{b}), \bar{b}).$$

Case 1: $a \in \mathcal{S}(2 \cdot 3^j, \bar{a})$.

Then $\mathcal{S}(3^j, \bar{a}a) \subseteq \mathcal{S}(3^{j+1}, \bar{a})$. Setting $b := \pi(a)$, we have $\pi : (\mathcal{S}(3^j, \bar{a}a), \bar{a}a) \cong (\mathcal{S}(3^j, \bar{b}b), \bar{b}b)$, hence $\bar{a}a \mapsto \bar{b}b \in I_j$.

Case 2: $a \notin \mathcal{S}(2 \cdot 3^j, \bar{a})$ (and hence, $\mathcal{S}(3^j, a) \cap \mathcal{S}(3^j, \bar{a}) = \emptyset$).

Let ι be the 3^j -ball type of a . By (*), $\mathcal{S}(2 \cdot 3^j, \bar{a})$ and $\mathcal{S}(2 \cdot 3^j, \bar{b})$ contain the same number of elements of 3^j -ball type ι which, by our assumption on the cardinality of balls, is $\leq \text{length}(\bar{a}) \cdot e \leq m \cdot e$. Therefore, by (i) or

⁶ Note that for finite \mathcal{A} and \mathcal{B} there is always such an e .

(ii), there must be an element $b \notin S(2 \cdot 3^j, \bar{b})$ with 3^j -ball type ι . Choose $\pi' : (\mathcal{S}(3^j, a), a) \cong (\mathcal{S}(3^j, b), b)$. Then the corresponding restriction of $\pi \cup \pi'$ is an isomorphism of $(\mathcal{S}(3^j, \bar{a}a), \bar{a}a)$ onto $(\mathcal{S}(3^j, \bar{b}b), \bar{b}b)$. \square

We give an application. Note that a graph \mathcal{G} is connected if each nonempty subset of G closed under the graph relation $E^{\mathcal{G}}$ contains already all elements of G , i.e., if \mathcal{G} is a model of the “second-order sentence”

$$\forall P((\exists xPx \wedge \forall x\forall y((Px \wedge Exy) \rightarrow Py)) \rightarrow \forall zPz).$$

We are going to show that the class of connected graphs is not axiomatizable by a second-order sentence of the form $\exists P_1 \dots \exists P_r \psi$, where P_1, \dots, P_r are unary and ψ is first-order (see section 3.1 for a precise definition of second-order logic).

For $l \geq 1$, let $\mathcal{D}_l = (D_l, E_l)$ be a digraph consisting of a cycle of length $l + 1$, e.g.

$$D_l := \{0, \dots, l\}, \quad E_l := \{(i, i + 1) \mid i < l\} \cup \{(l, 0)\}.$$

Lemma 2.4.2 *Suppose $\tau = \{E, P_1, \dots, P_r\}$ where P_1, \dots, P_r are unary, and let $m \geq 0$. Then there is an $l_0 \geq 1$ such that for any $l \geq l_0$ and any τ -structure of the form $(\mathcal{D}_l, P_1, \dots, P_r)$ there are $a, b \in D_l$ with disjoint and isomorphic 3^m -balls.*


Proof. For the structures under consideration any 3^m -ball contains exactly $2 \cdot 3^m + 1$ elements (note that P_1, \dots, P_r are unary and therefore do not influence the distances induced by the underlying digraphs). Let i be the number of possible isomorphism types of 3^m -balls. Then in a structure of cardinality $\geq l_0 := (i + 1)(2 \cdot 3^m + 1)$ there must be two points with disjoint 3^m -balls of the same isomorphism type. \square

Lemma 2.4.3 *Suppose $(\mathcal{D}_l, P_1, \dots, P_r)$ is a τ -structure (τ as in the preceding lemma) containing elements a and b with disjoint and isomorphic 3^m -balls. Denote by a_- and b_- the elements of D_l with $E_l a_- a$ and $E_l b_- b$, respectively (see the figure below). Let $(D_l, E'_l, P_1, \dots, P_r)$ be the structure obtained by splitting the cycle $(\mathcal{D}_l, P_1, \dots, P_r)$ into two cycles by removing the edges $(a_-, a), (b_-, b)$ and adding edges $(b_-, a), (a_-, b)$ instead; more formally:*

$$E'_l := (E_l \setminus \{(a_-, a), (b_-, b)\}) \cup \{(b_-, a), (a_-, b)\}.$$

Then $(\mathcal{D}_l, P_1, \dots, P_r) \cong_m (D_l, E'_l, P_1, \dots, P_r)$.

Proof. Immediate by Hanf’s Theorem, since both structures have the same number of 3^m -balls of any given isomorphism type. \square



Since a partial isomorphism between digraphs is a partial isomorphism of the associated graphs, we get from the two preceding lemmas:

Lemma 2.4.4 *For $\tau = \{E, P_1, \dots, P_r\}$ and $m \geq 0$ choose l_0 according to 2.4.2. Let $l \geq l_0$ and $(\mathcal{G}_l, P_1, \dots, P_r)$ be a τ -structure, where \mathcal{G}_l is the Gaifman graph $\mathcal{G}(\mathcal{D}_l)$ of \mathcal{D}_l , that is, \mathcal{G}_l is a cycle of length $l + 1$. Let \mathcal{G}'_l be the Gaifman graph $\mathcal{G}((\mathcal{D}_l, E'_l))$, where (\mathcal{D}_l, E'_l) is defined as in the preceding lemma. Then*

$$(\mathcal{G}_l, P_1, \dots, P_r) \equiv_m (\mathcal{G}'_l, P_1, \dots, P_r). \quad \square$$

We are now in a position to show:

Proposition 2.4.5 *The class of finite and connected graphs cannot be axiomatized by a formula of the form*

$$(*) \quad \exists P_1 \dots \exists P_r \psi,$$

where P_1, \dots, P_r are unary relation symbols and ψ is a first-order sentence over the vocabulary $\{E, P_1, \dots, P_r\}$.

Proof. Suppose that for the sentence $(*)$ and any finite graph \mathcal{G} , we have:

$$\mathcal{G} \text{ is connected iff for some } P_1, \dots, P_r \subseteq G : (\mathcal{G}, P_1, \dots, P_r) \models \psi.$$

For $m := \text{qr}(\psi)$ choose l_0 as in 2.4.2. Since \mathcal{G}_{l_0} is connected, there are P_1, \dots, P_r such that $(\mathcal{G}_{l_0}, P_1, \dots, P_r) \models \psi$. Then, $(\mathcal{G}'_{l_0}, P_1, \dots, P_r) \models \psi$ by 2.4.4, but \mathcal{G}'_{l_0} is not connected, a contradiction. \square

On the other side we have

Proposition 2.4.6 *The class of finite and connected graphs can be axiomatized by a formula of the form $\exists R \psi$, where R is binary and ψ is a first-order sentence over the vocabulary $\{E, R\}$.*

Proof. Let ψ be a sentence expressing that R is an irreflexive and transitive relation with a minimal element, and that Exy holds for any immediate R -successor y of x ; that is, ψ is the conjunction of

$$\begin{aligned} & \forall x \neg Rxx \wedge \forall x \forall y \forall z ((Rxy \wedge Ryz) \rightarrow Rxz), \\ & \exists x \forall y (x = y \vee Rxy), \\ & \forall x \forall y ((Rxy \wedge \forall z \neg (Rzx \wedge Rzy)) \rightarrow Exy). \end{aligned}$$

Let \mathcal{G} be a graph. Clearly, if \mathcal{G} is a model of $\exists R\psi$, say $(\mathcal{G}, R^A) \models \psi$, then for any element of \mathcal{G} there is a path connecting it with the minimal element; hence, \mathcal{G} is connected. Conversely, suppose \mathcal{G} is connected. Choose an arbitrary $a \in G$. For $n \in \mathbb{N}$ set $L_n := \{b \mid d(a, b) = n\}$ and take as R the transitive closure of $\{(b, c) \mid E^G bc \text{ and for some } n, b \in L_n \text{ and } c \in L_{n+1}\}$. \square

Exercise 2.4.7 Show the results corresponding to 2.4.5 and 2.4.6 for the class of finite acyclic digraphs (cf. 1.A1). (Hint: Denote by \mathcal{H}_l the digraph $(\{0, \dots, l\}, \{(i, i+1) \mid i < l\})$. For the nonaxiomatizability result use Hanf's Theorem similarly as above, cutting off in some $(\mathcal{H}_l, P_1, \dots, P_r)$ an interval with endpoints a and b for suitable a and b and forming a cycle out of it.) Moreover show that the class of finite acyclic digraphs can be axiomatized by a sentence $\forall P\psi$, where P is unary and ψ is a first-order sentence over $\{E, P\}$. \square

Exercise 2.4.8 Show that there is a formula $\psi(x, y)$ of the form $\exists P\varphi$, where P is unary and φ is a first-order formula over $\{E, P\}$, expressing in finite graphs that x and y are in the same connected component. In fact, as $\psi(x, y)$ one can take a formula expressing that $x = y$ or that there is a subset P (a “path from x to y ”) containing x and y such that both x and y have an edge to exactly one member of P and every other member of P has an edge to precisely two members of P . Conclude that $\forall x \forall y \exists P\varphi$ is not equivalent to a sentence of the form $\exists P_1 \dots \exists P_r \chi$ with unary P_1, \dots, P_r and first-order χ . \square

2.5 Gaifman's Theorem

Fix a relational τ . Let \mathcal{A} be a τ -structure. A subset M of A is l -scattered, if the distance (in the Gaifman graph $\mathcal{G}(\mathcal{A})$) between any two elements of M exceeds l . Given $r, n \geq 1$ and a τ -formula $\psi(x)$, it is easy to write down (see below) a first-order sentence asserting that there is a $2r$ -scattered subset M of cardinality at least n such that $\mathcal{S}(r, a) \models \psi[a]$ for all $a \in M$.⁷ Gaifman's Theorem states that every first-order sentence is logically equivalent to a boolean combination of such sentences. It thus is a further formulation of

⁷ Note that, due to $2r$ -scatteredness, the balls $\mathcal{S}(r, a)$ for $a \in M$ are pairwise disjoint.

the fact already present in Hanf's Theorem that first-order sentences only capture local properties of structures.

In order to give a more precise formulation we need some notations. First note that there is a τ -formula $\theta_n(x, y)$ such that for any τ -structure \mathcal{A} and $a, b \in A$,

$$\mathcal{A} \models \theta_n(x, y) [a, b] \quad \text{iff} \quad d(a, b) \leq n.$$

In fact, set $\theta_0(x, y) := x = y$ and (denoting by $a(R)$ the arity of R)

$$\begin{aligned} \theta_{n+1}(x, y) &:= \theta_n(x, y) \vee \exists z(\theta_n(x, z) \wedge \bigvee_{R \in \tau} \exists u_1 \dots \exists u_{a(R)} \\ &\quad (Ru_1 \dots u_{a(R)} \wedge \bigvee_{1 \leq i, j \leq a(R)} (u_i = z \wedge u_j = y))). \end{aligned}$$

From now on we write $d(x, y) \leq n$ for $\theta_n(x, y)$. For $\bar{x} = x_1 \dots x_m$ let

$$d(\bar{x}, y) \leq n := (d(x_1, y) \leq n \vee \dots \vee d(x_m, y) \leq n).$$

Let $k \in \mathbb{N}$. With every τ -formula $\varphi = \varphi(\bar{x}, \bar{y})$ we associate a formula $\varphi^{S(k, \bar{x})}(\bar{x}, \bar{y})$ such that for any τ -structure \mathcal{A} , $\bar{a} \in A$, and $\bar{b} \in S(k, \bar{a})$,

$$\mathcal{A} \models \varphi^{S(k, \bar{x})}[\bar{a}, \bar{b}] \quad \text{iff} \quad S(k, \bar{a}) \models \varphi[\bar{a}, \bar{b}].$$

To define $\varphi^{S(k, \bar{x})}$, first replace any bound occurrence in φ of a variable in \bar{x} by a new variable and then inductively relativize the quantifiers to $S(k, \bar{x})$, e.g.,

$$[\exists z \varphi]^{S(k, \bar{x})} := \exists z(d(\bar{x}, z) \leq k \wedge \varphi^{S(k, \bar{x})}).$$

Call a sentence *basic local* if it has the form

$$\exists x_1 \dots \exists x_n \bigwedge_{1 \leq i < j \leq n} (d(x_i, x_j) > 2r \wedge \psi^{S(r, x_i)}(x_i)),$$

where $\psi = \psi(x)$ is a first-order formula. Note that for $l \leq k$,

$$\models \varphi^{S(l, \bar{x})} \leftrightarrow [\varphi^{S(k, \bar{x})}]^{S(l, \bar{x})} \quad \text{and} \quad \models \varphi^{S(l, \bar{x})} \leftrightarrow [\varphi^{S(l, \bar{x})}]^{S(k, \bar{x})};$$

in particular, any sentence of the form

$$\exists x_1 \dots \exists x_n \bigwedge_{1 \leq i < j \leq n} (d(x_i, x_j) > 2r \wedge \psi^{S(l, x_i)}(x_i))$$

with $l \leq r$ is logically equivalent to a basic local sentence.

A *local* sentence is a boolean combination of basic local sentences.

Theorem 2.5.1 (Gaifman's Theorem) *Every first-order sentence is logically equivalent to a local sentence.*

By 1.0.7 it suffices to show

Lemma 2.5.2 *Suppose \mathcal{A} and \mathcal{B} satisfy the same basic local sentences. Then $\mathcal{A} \equiv \mathcal{B}$.*

Proof. We show that $\mathcal{A} \cong_m \mathcal{B}$ for $m \in \mathbb{N}$. The argument parallels that for Hanf's Theorem. There the sets I_j consisted of partial isomorphisms $\bar{a} \mapsto \bar{b}$ such that $\text{length}(\bar{a}) \leq m - j$ and

$$(\mathcal{S}(3^j, \bar{a}), \bar{a}) \cong (\mathcal{S}(3^j, \bar{b}), \bar{b}).$$

Here we replace \cong by $\equiv_{g(j)}$ and take balls of radius 7^j ; the values $g(0), g(1), \dots$ of the function g can be defined by induction: $g(j)$ only has to be greater than some values which one gets in the course of the proof.

So let I_j comprise all the partial isomorphisms $\bar{a} \mapsto \bar{b}$ from \mathcal{A} to \mathcal{B} such that $\text{length}(\bar{a}) \leq m - j$ and

$$(\mathcal{S}(7^j, \bar{a}), \bar{a}) \equiv_{g(j)} (\mathcal{S}(7^j, \bar{b}), \bar{b}).$$

Again, in case $\text{length}(\bar{a}) = 0$, we set $(\mathcal{S}(7^j, \bar{a}), \bar{a}) = \emptyset$ and agree that $\emptyset \equiv_k \emptyset$ for all k . In particular, $\emptyset \mapsto \emptyset \in I_m$. We show $(I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$. By symmetry, we can restrict ourselves to the forth property. Thus suppose $0 \leq j < m$, $a \in \mathcal{A}$, $\bar{a} \mapsto \bar{b} \in I_{j+1}$; hence

$$(1) \quad (\mathcal{S}(7^{j+1}, \bar{a}), \bar{a}) \equiv_{g(j+1)} (\mathcal{S}(7^{j+1}, \bar{b}), \bar{b}).$$

We introduce a useful abbreviation: For \bar{d} in a structure \mathcal{D} let

$$\psi_{\bar{d}}^j(\bar{x}) := \left[\varphi_{\mathcal{S}(7^j, \bar{d}), \bar{d}}^{g(j)}(\bar{x}) \right]^{\mathcal{S}(7^j, \bar{x})}$$

that is, $\psi_{\bar{d}}^j(\bar{x})$ expresses that $(\mathcal{S}(7^j, \bar{d}), \bar{d}) \equiv_{g(j)} (\mathcal{S}(7^j, \bar{x}), \bar{x})$ (recall that $\varphi_{\mathcal{D}, \bar{d}}^l$ denotes the l -isomorphism type of \bar{d} in \mathcal{D}).

Case 1. $a \in \mathcal{S}(2 \cdot 7^j, \bar{a})$.

Then

$$\mathcal{S}(7^{j+1}, \bar{a}) \models \exists z (d(\bar{a}, z) \leq 2 \cdot 7^j \wedge \psi_{\bar{a}a}^j(\bar{a}z)).$$

We assume that the quantifier rank of this formula is $\leq g(j+1)$ (this gives us a first condition on the value of $g(j+1)$). Hence by (1),

$$\mathcal{S}(7^{j+1}, \bar{b}) \models \exists z (d(\bar{b}, z) \leq 2 \cdot 7^j \wedge \psi_{\bar{b}a}^j(\bar{b}z)),$$

so that for some b we have

$$(\mathcal{S}(7^j, \bar{a}a), \bar{a}a) \equiv_{g(j)} (\mathcal{S}(7^j, \bar{b}b), \bar{b}b).$$

Therefore $\bar{a}a \mapsto \bar{b}b \in I_j$.

Case 2. $a \notin S(2 \cdot 7^j, \bar{a})$, i.e., $S(7^j, \bar{a}) \cap S(7^j, a) = \emptyset$.

For $s \geq 1$, the following formula $\delta_s(x_1, \dots, x_s)$ expresses that $\{x_1, \dots, x_s\}$ is a $4 \cdot 7^j$ -scattered set of elements whose 7^j -ball has the same $g(j)$ -isomorphism type as that of a :

$$\delta_s := \bigwedge_{1 \leq l < k \leq s} d(x_l, x_k) > 4 \cdot 7^j \wedge \bigwedge_{1 \leq l \leq s} \psi_a^j(x_l).$$

Compare the cardinalities (e and i below) of maximal $4 \cdot 7^j$ -scattered sets in $S(2 \cdot 7^j, \bar{a})$ and in \mathcal{A} , respectively, consisting of such elements. More precisely, let e and i be such that

$$(2) \quad S(7^{j+1}, \bar{a}) \models \exists x_1 \dots \exists x_e (\bigwedge_{1 \leq k \leq e} d(\bar{a}, x_k) \leq 2 \cdot 7^j \wedge \delta_e)$$

$$(3) \quad S(7^{j+1}, \bar{a}) \not\models \exists x_1 \dots \exists x_{e+1} (\bigwedge_{1 \leq k \leq e+1} d(\bar{a}, x_k) \leq 2 \cdot 7^j \wedge \delta_{e+1})$$

$$(4) \quad \mathcal{A} \models \exists x_1 \dots \exists x_i \delta_i, \quad \mathcal{A} \not\models \exists x_1 \dots \exists x_{i+1} \delta_{i+1};$$

if no such i exists, set $i = \infty$ (note that e is bounded by the length of \bar{a} (and hence by m), since any two elements of the same ball of radius $2 \cdot 7^j$ have a distance at most $4 \cdot 7^j$). Clearly, $e \leq i$. Moreover we claim that the corresponding numbers e and i determined in $S(7^{j+1}, \bar{b})$ and \mathcal{B} , respectively, are the same: Concerning \mathcal{B} , this holds since the sentences in (4) are basic local up to logical equivalence. Concerning $S(7^{j+1}, \bar{b})$ note that $(S(7^{j+1}, \bar{a}), \bar{a}) \equiv_{g(j+1)} (S(7^{j+1}, \bar{b}), \bar{b})$ and that $g(j+1)$ is greater than the quantifier rank of the sentences in (2) and (3) (this gives a second condition on the value of $g(j+1)$; recall that e is bounded by m).

Case 2.1. $e = i$.

Then all elements satisfying $\psi_a^j(x)$ have a distance

$$\leq 4 \cdot 7^j + 2 \cdot 7^j = 6 \cdot 7^j < 7^{j+1}$$

from \bar{a} (in fact, if one such element a' satisfies $d(\bar{a}, a') > 6 \cdot 7^j$ then a' together with e many witnesses for (2) show that $i \geq e + 1$). In particular, this holds for a . Since $a \notin S(2 \cdot 7^{j+1}, \bar{a})$,

$$S(7^{j+1}, \bar{a}) \models \exists z (2 \cdot 7^j < d(\bar{a}, z) \leq 6 \cdot 7^j \wedge \psi_a^j(z) \wedge \psi_a^j(\bar{a})).$$

Then, by (1),

$$S(7^{j+1}, \bar{b}) \models \exists z (2 \cdot 7^j < d(\bar{b}, z) \leq 6 \cdot 7^j \wedge \psi_a^j(z) \wedge \psi_a^j(\bar{b}))$$

(this gives us a third condition on the value of $g(j+1)$). Thus there is b with $2 \cdot 7^j < d(\bar{b}, b) \leq 6 \cdot 7^j$, and

$$(5) \quad (\mathcal{S}(7^j, a), a) \equiv_{g(j)} (\mathcal{S}(7^j, b), b).$$

Moreover,

$$(6) \quad (\mathcal{S}(7^j, \bar{a}), \bar{a}) \equiv_{g(j)} (\mathcal{S}(7^j, \bar{b}), \bar{b}).$$

Since the universes of the structures on the left sides of (5) and (6) are disjoint, and the same applies to the right sides, we obtain from 2.3.10(b) that

$$(\mathcal{S}(7^j, \bar{a}a), \bar{a}a) \equiv_{g(j)} (\mathcal{S}(7^j, \bar{b}b), \bar{b}b),$$

thus $\bar{a}a \mapsto \bar{b}b \in I_j$.

Case 2.2. $e < i$.

Then

$$\mathcal{B} \models \exists x_1 \dots \exists x_{e+1} \delta_{e+1}.$$

Hence there must be an element b in \mathcal{B} such that

$$\mathcal{S}(7^j, \bar{b}) \cap \mathcal{S}(7^j, b) = \emptyset \quad \text{and} \quad \mathcal{B} \models \psi_a^j(x)[b],$$

in particular, $(\mathcal{S}(7^j, a), a) \equiv_{g(j)} (\mathcal{S}(7^j, b), b)$. Now one can argue as at the end of the preceding case. \square

We close with an application of Gaifman's Theorem which will be used in Chapter 9. In the rest of this section all structures are assumed to be finite. Recall that τ is relational. Given τ -structures \mathcal{A} and \mathcal{B} , a mapping $h : A \rightarrow B$ is a *homomorphism* if for all $R \in \tau$ and $\bar{a} \in A$, $R\bar{a}$ implies $Rh(\bar{a})$. The homomorphism is said to be *strict* if, in addition, for all $R \in \tau$ and $\bar{a} \in A$ with $R^B h(\bar{a})$ there is $\bar{e} \in A$ such that $R^A \bar{e}$ and $h(\bar{e}) = h(\bar{a})$.

A sentence φ is *preserved under (strict) homomorphisms* if for all \mathcal{A}, \mathcal{B} and any (strict) homomorphism $h : A \rightarrow B$,

$$\mathcal{A} \models \varphi \quad \text{implies} \quad \mathcal{B} \models \varphi.$$

Exercise 2.5.3 Every existential positive sentence, that is, every sentence built up from atomic formulas with the connectives \wedge and \vee , and the quantifier \exists , is preserved under homomorphisms. \square

A model \mathcal{A} of a sentence φ is said to be *minimal* if no proper substructure is a model of φ , that is, if

$$\mathcal{B} \subseteq \mathcal{A} \text{ and } \mathcal{B} \models \varphi \quad \text{imply} \quad \mathcal{B} = \mathcal{A}.$$

Exercise 2.5.4 For any φ , every model of φ contains a minimal model, that is, if $\mathcal{B} \models \varphi$ then there is $\mathcal{A} \subseteq \mathcal{B}$ such that \mathcal{A} is a minimal model of φ (recall that we restrict ourselves to finite structures). \square

Theorem 2.5.5 *If φ is preserved under strict homomorphisms then there are l and m such that no minimal model of φ contains an l -scattered subset of cardinality m .*

Proof. By Gaifman's Theorem, φ is logically equivalent to a boolean combination of basic local sentences $\varphi_1, \dots, \varphi_k$. Suppose

$$\varphi_i = \exists x_1 \dots \exists x_{n_i} \left(\bigwedge_{1 \leq s < t \leq n_i} (d(x_s, x_t) > 2r_i \wedge \psi_i^{S(r_i, x_s)}(x_s)) \right).$$

Set $r := \max\{r_i \mid 1 \leq i \leq k\}$, $l := 2r$ and $m := 2^k + 1$. Let \mathcal{A} be a minimal model of φ . We show that \mathcal{A} contains no l -scattered subset of cardinality m . By contradiction, suppose that M is l -scattered and $\|M\| \geq m$. For $i = 1, \dots, k$ let $\rho_i(u)$ express “there is v such that $d(u, v) \leq r_i$ and $\psi_i^{S(r_i, v)}(v)$ ”. By choice of m there are $a, a' \in M$ such that $a \neq a'$ and for $i = 1, \dots, k$,

$$\mathcal{A} \models \rho_i[a] \text{ iff } \mathcal{A} \models \rho_i[a'].$$

Let \mathcal{B} be the substructure of \mathcal{A} with universe $A \setminus \{a\}$. Then $\mathcal{B} \not\models \varphi$, since \mathcal{A} is a minimal model of φ .

Set $n := \max\{n_i \mid i = 1, \dots, k\}$, $\mathcal{B}_n := \dot{\bigcup}_{j=1, \dots, n} \mathcal{B}$ (the disjoint union of n copies of \mathcal{B}), and $\mathcal{A}_n := \mathcal{A} \dot{\cup} \mathcal{B}_n$. The projection of \mathcal{B}_n to \mathcal{B} is a strict homomorphism, therefore, $\mathcal{B}_n \not\models \varphi$. Since the inclusion map of \mathcal{A} to \mathcal{A}_n is a strict homomorphism, too, and $\mathcal{A} \models \varphi$, we have $\mathcal{A}_n \models \varphi$. We obtain the desired contradiction, if we show that for $i = 1, \dots, k$,

$$\mathcal{A}_n \models \varphi_i \text{ iff } \mathcal{B}_n \models \varphi_i.$$

Fix i . Suppose first that $\mathcal{B}_n \models \varphi_i$. Then, for some $b \in B_n$, $S^{B_n}(r_i, b) \models \psi_i[b]$. View b as an element of \mathcal{B} . Then the B_n -part of \mathcal{A}_n contains n_i – even n – copies of the element b , which are pairwise at infinite distance. Since $S^{B_n}(r_i, b) \cong S^{A_n}(r_i, b)$, we obtain $\mathcal{A}_n \models \varphi_i$.

Assume now that $\mathcal{A}_n \models \varphi_i$. Choose $e \in A_n$ such that $S^{A_n}(r_i, e) \models \psi_i[e]$. If $a \notin S^A(r_i, e)$ then $S^{A_n}(r_i, e) \cong S^B(r_i, e)$, and we argue as above. Finally, assume that $a \in S^A(r_i, e)$. Then, $\mathcal{A} \models \rho_i[a]$, and hence, $\mathcal{A} \models \rho_i[a']$. Thus there is $e' \in A$ such that $d(e', a') \leq r$ and $S^A(r_i, e') \models \psi_i[e']$. Now $d(e', a) \geq d(a', a) - d(a', e') > l - r = 2r - r \geq r_i$. Therefore, $a \notin S^A(r_i, e')$, and hence $S^A(r_i, e') \cong S^B(r_i, e')$. Therefore, \mathcal{B}_n contains n_i copies of e' , and thus, $\mathcal{B}_n \models \varphi_i$. \square

Notes 2.5.6 As already mentioned in the introduction to this chapter, the algebraic characterization of m -equivalence is due to Fraïssé [43], its game-theoretic version to Ehrenfeucht [33]. The j -isomorphism types were introduced by Hintikka [85] in a different context. Theorem 2.4.1 is due to Hanf [79], Theorem 2.5.1 to Gaifman [45], and Theorem 2.5.5 to Ajtai and Gurevich [7]. Further references for the results in this chapter are [5, 78, 39]. The “local” character as it becomes apparent for first-order logic in the theorems of Hanf and Gaifman is studied in [83, 115, 116, 69].

3. More on Games

In this section we show that for some fragments and extensions of first-order logic there are corresponding variants of the Ehrenfeucht-Fraïssé method. First we deal with second-order logic, mainly with monadic second-order logic, then with infinitary languages, and finally with restrictions of these infinitary languages and of first-order language to fragments consisting of formulas which only contain a fixed finite number of variables.

3.1 Second-Order Logic

Second-order logic, SO, is an extension of first-order logic which allows to quantify over relations. In addition to the symbols of first-order logic, its alphabet contains, for each $n \geq 1$, countably many n -ary relation (or *predicate*) variables V_1^n, V_2^n, \dots . To denote relation variables we use letters X, Y, \dots . We define the set of second-order formulas of vocabulary τ to be the set generated by the rules for first-order formulas extended by:

- If X is n -ary and t_1, \dots, t_n are terms then $Xt_1 \dots t_n$ is a formula.
- If φ is a formula and X is a relation variable then $\exists X\varphi$ is a formula.

The free occurrence of a variable or of a relation variable in a second-order formula is defined in the obvious way and the notion of satisfaction is extended canonically. Then, given $\varphi = \varphi(x_1, \dots, x_n, Y_1, \dots, Y_k)$ with free (individual and relation) variables among $x_1, \dots, x_n, Y_1, \dots, Y_k$, a τ -structure \mathcal{A} , elements $a_1, \dots, a_n \in A$, and relations R_1, \dots, R_k over A of arities corresponding to Y_1, \dots, Y_k , respectively,

$$\mathcal{A} \models \varphi[a_1, \dots, a_n, R_1, \dots, R_k]$$

means that a_1, \dots, a_n together with R_1, \dots, R_k satisfy φ in \mathcal{A} .

For any τ the class $\text{EVEN}[\tau]$ of finite τ -structures of even cardinality is axiomatizable in second-order logic (but not in first-order logic, as we saw in 2.3.5). In fact, $\text{EVEN}[\tau] = \text{Mod}(\varphi)$, where φ is a sentence expressing “there is a binary relation which is an equivalence relation having only equivalence classes with exactly two elements”, e.g.,

$$\exists X(\forall xXxx \wedge \forall x\forall y(Xxy \rightarrow Xyx) \wedge \forall x\forall y\forall z((Xxy \wedge Xyz) \rightarrow Xxz) \\ \wedge \forall x\exists^=1y(Xxy \wedge y \neq x)).$$

We are mainly interested in the fragment MSO of second-order logic known as *monadic second-order logic*. In formulas of MSO only *unary* relation variables (“set variables”) are allowed. We write $\mathcal{A} \equiv_m^{\text{MSO}} \mathcal{B}$ if \mathcal{A} and \mathcal{B} satisfy the same monadic second-order sentences of quantifier rank $\leq m$ (the quantifier rank is the maximal number of nested first-order and second-order quantifiers).

As in first-order logic, \equiv_m^{MSO} can be characterized by an Ehrenfeucht-Fraïssé game, $\text{MSO-G}_m(\mathcal{A}, \mathcal{B})$. The rules are the same as in the first-order Ehrenfeucht-Fraïssé game, but now in every move the spoiler can decide whether to make a *point move* or a *set move*. The point moves are as the moves in the first-order case. In a set move the spoiler chooses a subset $P \subseteq A$ or $Q \subseteq B$, and then the duplicator answers by a subset $Q \subseteq B$ or $P \subseteq A$, respectively. After m moves, elements a_1, \dots, a_r and subsets P_1, \dots, P_s in A , and corresponding elements b_1, \dots, b_r and subsets Q_1, \dots, Q_s in B (with $m = r + s$) have been chosen. The duplicator wins if $\bar{a} \mapsto \bar{b} \in \text{Part}((\mathcal{A}, P_1, \dots, P_s), (\mathcal{B}, Q_1, \dots, Q_s))$.

Theorem 3.1.1 $\mathcal{A} \equiv_m^{\text{MSO}} \mathcal{B}$ iff the duplicator wins $\text{MSO-G}_m(\mathcal{A}, \mathcal{B})$.

The following exercise leads to a proof of this theorem (along the lines of the proof of the corresponding result 2.2.8). \square

Exercise 3.1.2 Given \mathcal{A} , $\bar{a} (= a_1 \dots a_r)$ in A , and $\bar{P} (= P_1 \dots P_s)$ a sequence of subsets of A , define the formulas $\psi_{\bar{a}, \bar{P}}^j$ similar to the j -isomorphism type $\varphi_{\bar{a}}^j$, but now taking into account also the second-order set quantifiers:

$$\psi_{\bar{a}, \bar{P}}^0 := \\ \bigwedge \{ \varphi(v_1, \dots, v_r, V_1, \dots, V_s) \mid \varphi \text{ atomic or negated atomic, } \mathcal{A} \models \varphi[\bar{a}, \bar{P}] \}$$

$$\psi_{\bar{a}, \bar{P}}^{j+1} := \\ \bigwedge_{a \in A} \exists v_{r+1} \psi_{\bar{a}a, \bar{P}}^j \wedge \forall v_{r+1} \bigvee_{a \in A} \psi_{\bar{a}a, \bar{P}}^j \wedge \bigwedge_{P \subseteq A} \exists V_{s+1} \psi_{\bar{a}, \bar{P}P}^j \wedge \forall V_{s+1} \bigvee_{P \subseteq A} \psi_{\bar{a}, \bar{P}P}^j.$$

Show the equivalence of

- (i) The duplicator wins $\text{MSO-G}_m((\mathcal{A}, \bar{P}, \bar{a}), (\mathcal{B}, \bar{Q}, \bar{b}))$;
- (ii) $\mathcal{B} \models \psi_{\bar{a}, \bar{P}}^m[\bar{b}, \bar{Q}]$;
- (iii) \bar{a}, \bar{P} satisfies in \mathcal{A} the same formulas of MSO of quantifier rank $\leq m$ as \bar{b}, \bar{Q} in \mathcal{B} . \square

As in first-order logic (compare 2.2.6) or by a direct proof, one easily gets the following result which we need later.

Proposition 3.1.3 For a fixed vocabulary and $m \in \mathbb{N}$, the relation \equiv_m^{MSO} is an equivalence relation with finitely many equivalence classes. \square

We give an application of the gametheoretic characterization.

Proposition 3.1.4 *The disjoint union and the ordered sum¹ preserve the relation \equiv_m^{MSO} , i.e., for relational τ we have:*

- (a) *If $\mathcal{A}_1 \equiv_m^{\text{MSO}} \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m^{\text{MSO}} \mathcal{B}_2$ then $\mathcal{A}_1 \dot{\cup} \mathcal{A}_2 \equiv_m^{\text{MSO}} \mathcal{B}_1 \dot{\cup} \mathcal{B}_2$.*
- (b) *If $\mathcal{A}_1 \equiv_m^{\text{MSO}} \mathcal{B}_1$ and $\mathcal{A}_2 \equiv_m^{\text{MSO}} \mathcal{B}_2$ then $\mathcal{A}_1 \triangleleft \mathcal{A}_2 \equiv_m^{\text{MSO}} \mathcal{B}_1 \triangleleft \mathcal{B}_2$.*

Proof. The proofs for (a) and (b) proceed in the same way. So let $* \in \{\dot{\cup}, \triangleleft\}$. Assume

$$\mathcal{A}_1 \equiv_m^{\text{MSO}} \mathcal{B}_1 \quad \text{and} \quad \mathcal{A}_2 \equiv_m^{\text{MSO}} \mathcal{B}_2.$$

By hypothesis and the last theorem there are winning strategies S_1 and S_2 for the duplicator in the games $\text{MSO-G}_m(\mathcal{A}_1, \mathcal{B}_1)$ and $\text{MSO-G}_m(\mathcal{A}_2, \mathcal{B}_2)$, respectively. Then the following represents a winning strategy for the duplicator in $\text{MSO-G}_m(\mathcal{A}_1 * \mathcal{A}_2, \mathcal{B}_1 * \mathcal{B}_2)$ (when describing it we use moves of plays in $\text{MSO-G}_m(\mathcal{A}_1, \mathcal{B}_1)$ and $\text{MSO-G}_m(\mathcal{A}_2, \mathcal{B}_2)$). Suppose first that the i -th move of the spoiler is a point move where he selects, say, $a \in \mathcal{A}_1 * \mathcal{A}_2$. Then the duplicator gets his answer by applying S_1 if $a \in \mathcal{A}_1$, and S_2 if $a \in \mathcal{A}_2$. Now assume that the spoiler selects, say, $P \subseteq \mathcal{A}_1 \cup \mathcal{A}_2$. Set $P_1 := P \cap \mathcal{A}_1$ and $P_2 := P \cap \mathcal{A}_2$. Let Q_1 and Q_2 be the selections of the duplicator according to S_1 and S_2 , respectively. Then, in the game $\text{MSO-G}_m(\mathcal{A}_1 * \mathcal{A}_2, \mathcal{B}_1 * \mathcal{B}_2)$, the duplicator chooses $Q_1 \cup Q_2$. \square

An easy induction using equivalences such as

$$\models \neg \exists X \varphi \leftrightarrow \forall X \neg \varphi, \quad \models (\varphi \vee \forall Y \psi) \leftrightarrow \forall Y (\varphi \vee \psi) \quad \text{if } Y \text{ is not free in } \varphi,$$

shows that each (M)SO-formula is logically equivalent to an (M)SO-formula in prenex normal form, that is, to a formula of the form

$$Q_1 \alpha_1 \dots Q_s \alpha_s \psi,$$


where $Q_1, \dots, Q_s \in \{\forall, \exists\}$, and where $\alpha_1, \dots, \alpha_s$ are first-order or second-order variables and ψ is quantifier-free. Moreover, since

$$\begin{aligned} \models \exists x Q_1 \alpha_1 \dots Q_s \alpha_s \psi &\leftrightarrow \exists X Q_1 \alpha_1 \dots Q_s \alpha_s (\exists^{-1} x X x \wedge \forall x (X x \rightarrow \psi)), \\ \models \forall x Q_1 \alpha_1 \dots Q_s \alpha_s \psi &\leftrightarrow \forall X Q_1 \alpha_1 \dots Q_s \alpha_s (\exists^{-1} x X x \rightarrow \forall x (X x \rightarrow \psi)), \end{aligned}$$

every (M)SO-formula is logically equivalent to one in prenex normal form in which each second-order quantifier precedes all first-order quantifiers. Such a formula is called $(\text{M})\Sigma_n^1$, if the string of second-order quantifiers consists of n consecutive blocks, where in each block all quantifiers are of the same type (i.e., all universal or all existential), adjacent blocks contain quantifiers of different type, and the first block is existential. $(\text{M})\Pi_n^1$ -formulas are defined in the same way, but now we require that the first block consists of universal quantifiers.

¹ Compare 1.A3 for the definition of disjoint union and ordered sum.

In particular, $\exists X \exists Y \forall Z \forall x \exists y \varphi$ with quantifier-free φ is a Σ_2^1 -formula. Clearly, the negation of a Σ_n^1 -formula is logically equivalent to a Π_n^1 -formula, and the negation of a Π_n^1 -formula to a Σ_n^1 -formula. Moreover, denoting by Δ_n^1 the set of formulas that are logically equivalent to both a Σ_n^1 -formula and a Π_n^1 -formula, we have up to logical equivalence



which can easily be verified by adding dummy variables. The same inclusions hold for the monadic classes.

It can be shown that for arbitrary models all the inclusions above are proper (this also holds for MSO). The question to what extent the hierarchies are proper in the finite is related to important questions of complexity theory (cf. Chapter 7). We have seen in section 2.4 that the class of finite, connected graphs is $M\Pi_1^1$ -axiomatizable but not $M\Sigma_1^1$ -axiomatizable. Thus, in the finite, $M\Sigma_1^1 \neq M\Pi_1^1$.

3.2 Infinitary Logic: The Logics $L_{\infty\omega}$ and $L_{\omega_1\omega}$

In recent years even *infinitary* logics have turned out to be relevant in the context of finite model theory. In fact, they contain some of the logics important in descriptive complexity theory (see Chapter 7), but – in contrast to them – their expressive power allows a manageable characterization in terms of games, similar to that of first-order logic. In this and the next section we introduce these logics and deal with their gametheoretic aspects.

The infinitary logics $L_{\infty\omega}$ and $L_{\omega_1\omega}$ allow arbitrary and countable disjunctions (and hence conjunctions), respectively. More formally: Let τ be a vocabulary. The class of $L_{\infty\omega}$ -formulas over τ is given by the following clauses:

- it contains all atomic first-order formulas over τ
- if φ is a formula then so is $\neg\varphi$
- if φ is a formula and x a variable then $\exists x\varphi$ is a formula
- if Ψ is a set of formulas then $\bigvee \Psi$ is a formula.

For $L_{\omega_1\omega}$ we replace the last clause by

- if Ψ is a *countable* set of formulas then $\bigvee \Psi$ is a formula.

The semantics is a direct extension of the semantics of first-order logic with $\bigvee \Psi$ being interpreted as the disjunction over all formulas in Ψ ; hence, neglecting the interpretation of the free variables,

$$\mathcal{A} \models \bigvee \Psi \quad \text{iff} \quad \text{for some } \psi \in \Psi, \mathcal{A} \models \psi.$$

We set $\bigwedge \Psi := \neg \bigvee \{\neg\psi \mid \psi \in \Psi\}$. Then $\bigwedge \Psi$ is interpreted as the conjunction over all formulas in Ψ . By identifying $(\varphi \vee \psi)$ with $\bigvee \{\varphi, \psi\}$ we see that $L_{\infty\omega}$ and $L_{\omega_1\omega}$ are extensions of first-order logic.

Examples 3.2.1 (a) For any τ , the models of the $L_{\omega_1\omega}$ -sentence

$$\bigvee \{\varphi_{=n} \mid n \geq 1\},$$

where $\varphi_{=n}$ is a first-order sentence expressing that the universe has cardinality n , are the finite τ -structures. The $L_{\omega_1\omega}$ -sentence $\bigvee \{\varphi_{=2n} \mid n \geq 1\}$ axiomatizes the class $\text{EVEN}[\tau]$. Similarly, if M is any nonempty set of positive natural numbers, then the class of models of the $L_{\omega_1\omega}$ -sentence $\bigvee \{\varphi_{=k} \mid k \in M\}$ corresponds to the query “ $\|A\| \in M?$ ”. In particular, we see that nonrecursive queries are $L_{\omega_1\omega}$ -definable. We get even more:

(b) Any class of finite structures is axiomatizable in $L_{\omega_1\omega}$. In fact, let K be a class of finite structures. Choose, using 2.1.2, a set Φ of first-order sentences such that $K = \text{Mod}(\Phi)$. Then $K = \text{Mod}(\varphi)$ for the $L_{\omega_1\omega}$ -sentence $\varphi := \bigwedge \Phi$.

(c) “Connectivity” is a property of graphs expressible in $L_{\omega_1\omega}$ by

$$\forall x \forall y (\neg x = y \rightarrow \bigvee \{\varphi_n(x, y) \mid n \geq 1\}),$$

where $\varphi_n(x, y)$ is a first-order formula saying that there is a path from x to y of length n ,

$$\varphi_n(x, y) := \exists z_0 \dots \exists z_n (z_0 = x \wedge z_n = y \wedge E z_0 z_1 \wedge \dots \wedge E z_{n-1} z_n). \quad \square$$

$L_{\infty\omega}$ -sentences are $L_{\infty\omega}$ -formulas without free variables. While $L_{\infty\omega}$ -formulas may have infinitely many free variables ($\bigvee \{\neg v_i = v_j \mid 1 \leq i < j\}$ is an example), it can easily be seen that subformulas of $L_{\infty\omega}$ -sentences only have finitely many free variables. In the following we restrict ourselves to $L_{\infty\omega}$ -formulas with only finitely many free variables.

The next two propositions – and in particular the proof of the first one – show that we need not more than $L_{\omega_1\omega}$ in the context of finite model theory, although it has become popular – and we follow this tradition – to mainly consider $L_{\infty\omega}$.

Proposition 3.2.2 (a) *In the finite, every $L_{\infty\omega}$ -formula $\varphi(\bar{x})$ is equivalent to an $L_{\omega_1\omega}$ -formula $\psi(\bar{x})$.*

(b) *Assume \mathcal{A} and \mathcal{B} are finite. For every $L_{\infty\omega}$ -formula $\varphi(\bar{x})$ there is an FO-formula $\psi(\bar{x})$ such that*

$$\mathcal{A} \models \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \psi(\bar{x})) \quad \text{and} \quad \mathcal{B} \models \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \psi(\bar{x})).$$

In both cases, (a) and (b), the formula ψ can be chosen such that $\text{free}(\psi) \subseteq \text{free}(\varphi)$ and that every variable occurring in ψ (free or bound) occurs in φ .

Proof. The proof of (a) and (b) is by induction over the rules for $L_{\infty\omega}$ -formulas. The translation procedure preserves the “structure” of formulas and only replaces infinitary disjunctions by countable ones (by finite ones in (b)). In the main step suppose that

$$\varphi(\bar{x}) = \bigvee \{\varphi_i(\bar{x}) \mid i \in I\}$$

is an $L_{\infty\omega}$ -formula. For each finite structure \mathcal{C} with universe $\{1, 2, \dots, \|\mathcal{C}\|\}$ (for $\mathcal{C} \in \{\mathcal{A}, \mathcal{B}\}$ in (b)) and each $\bar{c} \in \mathcal{C}$, if there exists an $i \in I$ such that $\mathcal{C} \models \varphi_i[\bar{c}]$, choose such an i . Let I_0 be the set of i 's chosen in this way. Then I_0 is countable (finite in (b)) and $\bigvee \{\varphi_i(\bar{x}) \mid i \in I\}$ and $\bigvee \{\varphi_i(\bar{x}) \mid i \in I_0\}$ are equivalent in the finite (in \mathcal{A} and \mathcal{B}). \square

Since every finite structure can be characterized in first-order logic, we obtain the following improvement of part (a) of the preceding proposition (cf. 3.2.1(b)):

Proposition 3.2.3 *In the finite, every $L_{\infty\omega}$ -formula $\varphi(\bar{x})$ is equivalent to a countable disjunction – and hence to a countable conjunction – of first-order formulas. In fact, in the finite, $\varphi(\bar{x})$ is equivalent to*

$$\bigvee \{\varphi_{\mathcal{A}, \bar{a}}^{\|\mathcal{A}\|+1}(\bar{x}) \mid \mathcal{A} \text{ finite, } \bar{a} \in \mathcal{A}, \mathcal{A} \models \varphi[\bar{a}]\}.$$

Proof. For simplicity we restrict ourselves to sentences. Let \mathcal{B} be a finite structure. If $\mathcal{B} \models \varphi$ then $\varphi_{\mathcal{B}}^{\|\mathcal{B}\|+1}$ is a member of the disjunction which, therefore, is satisfied by \mathcal{B} . Conversely, if \mathcal{B} satisfies the disjunction, then for some finite \mathcal{A} with $\mathcal{A} \models \varphi$ we have $\mathcal{B} \models \varphi_{\mathcal{A}}^{\|\mathcal{A}\|+1}$. Therefore by 2.2.10, $\mathcal{A} \cong \mathcal{B}$; hence, $\mathcal{B} \models \varphi$. \square

We say that \mathcal{A} and \mathcal{B} are $L_{\infty\omega}$ -equivalent, $\mathcal{A} \equiv^{L_{\infty\omega}} \mathcal{B}$, if \mathcal{A} and \mathcal{B} satisfy the same $L_{\infty\omega}$ -sentences. In order to characterize $\equiv^{L_{\infty\omega}}$, one needs Ehrenfeucht games with *infinitely* many moves.

Definition 3.2.4 Let \mathcal{A} and \mathcal{B} be structures, $\bar{a} \in A^s$, and $\bar{b} \in B^s$. The game $G_{\infty}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ is the same as the game $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ up to the fact that now each player has to make infinitely many moves. Thus, in the course of a play of $G_{\infty}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, elements e_1, e_2, \dots in \mathcal{A} and f_1, f_2, \dots in \mathcal{B} are chosen. The duplicator *wins the play* if $\bar{a}e_1 \dots e_i \mapsto \bar{b}f_1 \dots f_i \in \text{Part}(\mathcal{A}, \mathcal{B})$ for *all* i , and the spoiler wins if $\bar{a}e_1 \dots e_i \mapsto \bar{b}f_1 \dots f_i \notin \text{Part}(\mathcal{A}, \mathcal{B})$ for *some* i . The duplicator *wins* $G_{\infty}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ if he has a winning strategy. \square

Immediately from the definition we get:

Lemma 3.2.5 *Suppose that the duplicator wins $G_{\infty}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. Then*

- (a) $\bar{a} \mapsto \bar{b} \in \text{Part}(\mathcal{A}, \mathcal{B})$.
- (b) For $a \in A$ there is $b \in B$ such that the duplicator wins $G_{\infty}(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$.

(c) For $b \in B$ there is $a \in A$ such that the duplicator wins $G_\infty(\mathcal{A}, \bar{a}a, \mathcal{B}, \bar{b}b)$. \square

Since we want to give the Ehrenfeucht-Fraïssé type characterization of $\equiv^{L_{\infty\omega}}$ in its different forms we define:

Definition 3.2.6 (a) \mathcal{A} and \mathcal{B} are said to be *partially isomorphic*, written $\mathcal{A} \cong_{\text{part}} \mathcal{B}$, if there is a nonempty set I of partial isomorphisms from \mathcal{A} to \mathcal{B} with the back and forth properties (that is, for every $p \in I$ and every $a \in A$ ($b \in B$) there is $q \in I$ with $q \supseteq p$ and $a \in \text{do}(q)$ ($b \in \text{rg}(q)$)). We then write $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$.

(b) $W_\infty(\mathcal{A}, \mathcal{B}) :=$

$$\{\bar{a} \mapsto \bar{b} \mid s \in \mathbb{N}, \bar{a} \in A^s, \bar{b} \in B^s, \text{ the duplicator wins } G_\infty(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})\}$$

is the set of winning positions for the duplicator. \square

$L_{\infty\omega}$ -equivalence and the notions just introduced are intimately related:

Theorem 3.2.7 For structures \mathcal{A} and \mathcal{B} , $\bar{a} \in A^s$, and $\bar{b} \in B^s$ the following are equivalent:

- (i) The duplicator wins $G_\infty(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.
- (ii) $\bar{a} \mapsto \bar{b} \in W_\infty(\mathcal{A}, \mathcal{B})$ and $W_\infty(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}} \mathcal{B}$.
- (iii) There is a set I with $\bar{a} \mapsto \bar{b} \in I$ such that $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$.
- (iv) \bar{a} and \bar{b} satisfy the same formulas of $L_{\infty\omega}$ in \mathcal{A} and \mathcal{B} , respectively, that is, if $\varphi(x_1, \dots, x_s)$ is a formula of $L_{\infty\omega}$, then

$$\mathcal{A} \models \varphi[\bar{a}] \quad \text{iff} \quad \mathcal{B} \models \varphi[\bar{b}].$$

Proof. For (i) \Rightarrow (ii) see the preceding lemma; clearly (ii) implies (iii). For (iii) \Rightarrow (i) note that a set I with $\bar{a} \mapsto \bar{b} \in I$ and $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$ can be viewed as a winning strategy for the duplicator for the game $G_\infty(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. Hence, it suffices to show the equivalence of (iii) and (iv). Let I be as in (iii). By (transfinite) induction on the quantifier rank of the $L_{\infty\omega}$ -formula $\varphi(x_1, \dots, x_s)$ we prove

$$(*) \quad \mathcal{A} \models \varphi[\bar{e}] \quad \text{iff} \quad \mathcal{B} \models \varphi[\bar{f}]$$

for any $e_1 \dots e_s \mapsto f_1 \dots f_s \in I$. The case of quantifier rank 0 is handled by part (c) of 2.2.2. For arbitrary quantifier rank note that the class of formulas satisfying (*) contains the atomic formulas and is closed under \neg and \bigvee . Suppose that $\varphi(x_1, \dots, x_s) = \exists y \psi(x_1, \dots, x_s, y)$. Assume, for example, that $\mathcal{A} \models \varphi[e_1, \dots, e_s]$. Then there is $a \in A$ such that $\mathcal{A} \models \psi[e_1, \dots, e_s, a]$. By the forth property of I , there is $b \in B$ such that $e_1 \dots e_s a \mapsto f_1 \dots f_s b \in I$. Since $\text{qr}(\psi) < \text{qr}(\varphi)$, the induction hypothesis yields $\mathcal{B} \models \psi[f_1, \dots, f_s, b]$; hence, $\mathcal{B} \models \varphi[f_1, \dots, f_s]$.

Suppose now that (iv) holds, and let I be the set of all partial isomorphisms $e_1 \dots e_r \mapsto f_1 \dots f_r$ (with $r \geq 0$) such that for all $L_{\infty\omega}$ -formulas $\varphi(x_1, \dots, x_r)$,

$$\mathcal{A} \models \varphi[\bar{e}] \quad \text{iff} \quad \mathcal{B} \models \varphi[\bar{f}].$$

By (iv), $\bar{a} \mapsto \bar{b} \in I$. We show that I has the back and forth properties. So, let $e_1 \dots e_r \mapsto f_1 \dots f_r \in I$ and $a \in I$. For each $b \in B$, if there is a formula $\varphi(x_1, \dots, x_r, x)$ of $L_{\infty\omega}$ such that

$$\mathcal{A} \models \varphi(\bar{x}, x)[\bar{e}a] \quad \text{and} \quad \mathcal{B} \models \neg\varphi(\bar{x}, x)[\bar{f}b],$$

let $\varphi_b(\bar{x}, x)$ be such a formula; otherwise, set $\varphi_b(\bar{x}, x) := x = x$. Since $\mathcal{A} \models \exists x \bigwedge \{\varphi_b \mid b \in B\}[\bar{e}]$, we have $\mathcal{B} \models \exists x \bigwedge \{\varphi_b \mid b \in B\}[\bar{f}]$. Hence, there is $b' \in B$ such that $\mathcal{B} \models \bigwedge \{\varphi_b \mid b \in B\}[\bar{f}b']$. Using the definition of $\varphi_{b'}$, one easily sees that $\bar{e}a$ and $\bar{f}b'$ satisfy the same formulas of $L_{\infty\omega}$ in \mathcal{A} and \mathcal{B} , respectively, and hence, $\bar{e}a \mapsto \bar{f}b' \in I$. The back property is proven similarly. \square

Since $W_\infty(\mathcal{A}, \mathcal{B}) \neq \emptyset$ iff $\emptyset \mapsto \emptyset \in W_\infty(\mathcal{A}, \mathcal{B})$ (cf. 3.2.7), we obtain the

Corollary 3.2.8 *For \mathcal{A} and \mathcal{B} the following are equivalent.*

- (i) *The duplicator wins $G_\infty(\mathcal{A}, \mathcal{B})$.*
- (ii) $W_\infty(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}} \mathcal{B}$.
- (iii) $\mathcal{A} \cong_{\text{part}} \mathcal{B}$.
- (iv) $\mathcal{A} \cong^{L_{\infty\omega}} \mathcal{B}$. \square

To give some applications we first show:

Lemma 3.2.9 *Let A and B be countable.*

- (a) *If $\mathcal{A} \cong_{\text{part}} \mathcal{B}$ then $\mathcal{A} \cong \mathcal{B}$.*
- (b) *If $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$ and $p_0 \in I$ then p_0 can be extended to an isomorphism from \mathcal{A} onto \mathcal{B} .*

Proof. Let $A = \{a_1, a_2, \dots\}$ and $B = \{b_1, b_2, \dots\}$. It suffices to show (b). Suppose $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$ and $p_0 \in I$. By repeated application of the back and forth properties, one obtains p_1, p_2, \dots in I such that $p_0 \subseteq p_1 \subseteq \dots$ and such that $a_1 \in \text{do}(p_1), b_1 \in \text{rg}(p_2), a_2 \in \text{do}(p_3), \dots$. Then $\bigcup_{n \geq 0} p_n$ is an isomorphism from \mathcal{A} onto \mathcal{B} . \square

Corollary 3.2.10 *If \mathcal{A} and \mathcal{B} are countable and $L_{\infty\omega}$ -equivalent then they are isomorphic.*

Proof. The claim follows immediately from 3.2.8 and 3.2.9(a). \square

The following example will be of importance in the next chapter.

Example 3.2.11 Let τ be relational. For $r \geq 0$ let Δ_{r+1} be the set

$$\Delta_{r+1} := \{\varphi(v_1, \dots, v_r, v_{r+1}) \mid \varphi \text{ has the form } R\bar{x}, \text{ where } R \in \tau \text{ and} \\ \text{where } v_{r+1} \text{ occurs in } \bar{x}\}.$$

For a subset Φ of Δ_{r+1} the sentence $\chi_\Phi :=$ ²

$$\forall v_1 \dots \forall v_r \left(\bigwedge_{1 \leq i < j \leq r} v_i \neq v_j \rightarrow \exists v_{r+1} \left(\bigwedge_{1 \leq i \leq r} v_i \neq v_{r+1} \wedge \bigwedge_{\varphi \in \Phi} \varphi \wedge \bigwedge_{\varphi \in \Phi^c} \neg \varphi \right) \right)$$

where $\Phi^c := \Delta_{r+1} \setminus \Phi$, is called an *extension axiom*, more precisely, an $(r+1)$ -*extension axiom*. The set T_{rand} of all extension axioms gives the *random structure theory*.³ Clearly,

– every model of T_{rand} is infinite.

The extension axioms in T_{rand} guarantee for any models \mathcal{A} and \mathcal{B} of T_{rand} that the set

$$I := \{ \bar{a} \mapsto \bar{b} \mid \bar{a} \in A, \bar{b} \in B, \text{ and } \varphi_{\mathcal{A}, \bar{a}}^0 = \varphi_{\mathcal{B}, \bar{b}}^0 \}$$

has the back and forth properties. For a proof let $\bar{a} \mapsto \bar{b} \in I$, where $\bar{a} = a_1 \dots a_r$ and a_1, \dots, a_r can be assumed to be distinct, and let, say, a_{r+1} be in $A \setminus \{a_1, \dots, a_r\}$. Set $\Phi := \{ \varphi(v_1, \dots, v_{r+1}) \mid \varphi \in \Delta_{r+1}, \mathcal{A} \models \varphi[\bar{a}a_{r+1}] \}$. As $\mathcal{B} \models \chi_\Phi$, there is an element $b_{r+1} \in B$ such that $\varphi_{\mathcal{B}, \bar{b}b_{r+1}}^0 = \varphi_{\mathcal{A}, \bar{a}a_{r+1}}^0$, that is, $\bar{a}a_{r+1} \mapsto \bar{b}b_{r+1} \in I$. Moreover, since τ is relational, the empty partial isomorphism is in I . Hence, $I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$. By 3.2.8 we get:

– Any two models of T_{rand} are $L_{\infty\omega}$ -equivalent and thus, for each $L_{\infty\omega}$ -sentence φ , $T_{\text{rand}} \models \varphi$ or $T_{\text{rand}} \models \neg \varphi$.

Finally we show

– T_{rand} has a countable model and hence by 3.2.10, an (up to isomorphism) unique countable model \mathcal{R} , the so-called *infinite random structure*.

In fact, let $(\alpha_n)_{n \geq 0}$ be an enumeration of all pairs (\bar{m}, χ) , where \bar{m} is a tuple of distinct natural numbers and χ is an $(r+1)$ -extension axiom, where $r := \text{length}(\bar{m})$; moreover suppose that for $\alpha_n = (\bar{m}, \chi)$ all entries of \bar{m} are not greater than n . By induction on n we define structures \mathcal{A}_n with

$$A_n = \{0, \dots, n\} \text{ and } A_0 \subseteq A_1 \subseteq A_2 \subseteq \dots$$

such that $\mathcal{A} := \bigcup_{n \geq 0} \mathcal{A}_n$ is a model of T_{rand} : Let $\mathcal{A}_0 = (A_0, (\emptyset)_{R \in \tau})$ (each relation symbol is interpreted by the empty set). Suppose \mathcal{A}_n has been defined and $\alpha_n = (m_1, \dots, m_r, \chi)$ with $\chi = \chi_\Phi$. Define \mathcal{A}_{n+1} with universe A_{n+1} such that $\mathcal{A}_n \subseteq \mathcal{A}_{n+1}$ and such that for $\varphi \in \Delta_{r+1}$

$$\mathcal{A}_{n+1} \models \varphi[m_1, \dots, m_r, n+1] \quad \text{iff} \quad \varphi \in \Phi$$

(note that v_{r+1} occurs in every formula of Δ_{r+1}). This ensures that $\mathcal{A} := \bigcup_{n \geq 0} \mathcal{A}_n$ is a model of χ . \square

² In case $r = 0$ this sentence reduces to $\exists v_1 (\bigwedge_{\varphi \in \Phi} \varphi \wedge \bigwedge_{\varphi \in \Phi^c} \neg \varphi)$.

³ The name will become clear in Chapter 4.

We close with a further application of 3.2.8. For a sentence φ let the *spectrum* $\text{Spec}(\varphi)$ of φ be the set

$$\text{Spec}(\varphi) := \{m \geq 1 \mid \text{there is } \mathcal{A} \models \varphi \text{ with } \|\mathcal{A}\| = m\}.$$

Proposition 3.2.12 *For any first-order sentence φ at least one of $\text{Spec}(\varphi)$ or $\text{Spec}(\neg\varphi)$ is cofinite (i.e., for some n_0 , $\{n \mid n_0 \leq n\} \subseteq \text{Spec}(\varphi)$ or $\{n \mid n_0 \leq n\} \subseteq \text{Spec}(\neg\varphi)$).*

Proof. Let Φ be the set of the sentences

- (i) $\varphi_{\geq m}$ for $m \geq 1$
- (ii) $\forall \bar{x} R\bar{x}$ for $R \in \tau$
- (iii) $c = d$ for $c, d \in \tau$.

Clearly, Φ is satisfiable and any two models \mathcal{A} and \mathcal{B} of Φ are partially isomorphic via $I := \{p \in \text{Part}(\mathcal{A}, \mathcal{B}) \mid \text{do}(p) \text{ finite}\}$, and hence by 3.2.8, elementarily equivalent. Therefore, given a first-order sentence φ , we have $\Phi \models \varphi$ or $\Phi \models \neg\varphi$, say, $\Phi \models \varphi$. By the compactness theorem there is a finite $\Phi_0 \subseteq \Phi$ such that $\Phi_0 \models \varphi$. Let n_0 be larger than any m such that $\varphi_{\geq m}$ is in Φ_0 . Then Φ_0 and hence, φ has a model of cardinality n for each $n \geq n_0$. \square

Exercise 3.2.13 Prove the previous result for “vocabularies” containing function symbols. Hint: Compare the passage from such a vocabulary to a function-free one given in 1.D. For any new relation symbol F corresponding to some $f \in \tau$, in (ii) above take the sentence $\forall \bar{x} (\exists^=1 x F\bar{x}x \wedge F\bar{x}x_1)$ instead of $\forall \bar{x} x F\bar{x}x$. \square

Exercise 3.2.14 For any \mathcal{A} and \mathcal{B} we have

$$W_0(\mathcal{A}, \mathcal{B}) \supseteq \dots \supseteq W_m(\mathcal{A}, \mathcal{B}) \supseteq \dots \supseteq W_\infty(\mathcal{A}, \mathcal{B}).$$

Show that in case \mathcal{A} or \mathcal{B} is finite there is an $m_0 \leq 1 + \min\{\|\mathcal{A}\|, \|\mathcal{B}\|\}$ such that

$$W_0(\mathcal{A}, \mathcal{B}) \supset \dots \supset W_{m_0}(\mathcal{A}, \mathcal{B}) = W_\infty(\mathcal{A}, \mathcal{B}). \quad \square$$

3.3 The Logics FO^s and $\text{L}_{\infty\omega}^s$

In first-order logic FO, every finite structure \mathcal{A} can be characterized up to isomorphism by a first-order sentence $\varphi_{\mathcal{A}}$ which, in general, needs $\|\mathcal{A}\| + 1$ variables (compare, for example, the proof of 2.1.1). Hence, an arbitrary class K of finite structures can be axiomatized in $\text{L}_{\infty\omega}$ by the sentence

$$\bigvee \{\varphi_{\mathcal{A}} \mid \mathcal{A} \in K\}$$

which, in general, contains infinitely many variables. Since every class of finite structures is axiomatizable in it, $\text{L}_{\infty\omega}$ is too powerful in the finite to

yield new general principles. The observation above motivates the restriction to formulas of $L_{\infty\omega}$ containing only finitely many variables. In the present section we shall see that this approach really leads to logics which are interesting from the perspective of finite model theory.

Fix $s \geq 1$. We denote by $L_{\infty\omega}^s$ and FO^s the fragments of $L_{\infty\omega}$ and FO , respectively, containing only formulas, whose *free and bound* variables are among v_1, \dots, v_s . Moreover, we set

$$L_{\infty\omega}^\omega := \bigcup_{s \geq 1} L_{\infty\omega}^s.$$

Whereas $\text{FO} = \bigcup_{s \geq 1} \text{FO}^s$, the formula $\bigvee \{\varphi = n \mid n \geq 1\}$ belongs to $L_{\infty\omega}$ but not to $L_{\infty\omega}^\omega$.

Let $x = v_1, y = v_2$, and $z = v_3$ in examples.

Examples 3.3.1 (a) Let $\tau = \{<\}$. There are FO^2 -formulas $\psi_n(x)$ and χ_n such that for orderings \mathcal{A} , $a \in A$, and $n \geq 0$,

$$\mathcal{A} \models \psi_n[a] \quad \text{iff} \quad a \text{ is the } n\text{-th element of } <^{\mathcal{A}},$$

and in case $n \geq 1$,

$$\mathcal{A} \models \chi_n \quad \text{iff} \quad ||A|| = n.$$

In fact, define inductively

$$\psi_0(x) := \forall y \neg y < x, \quad \psi_{n+1}(x) := \forall y (y < x \leftrightarrow \bigvee_{i \leq n} \exists x (x = y \wedge \psi_i(x))),$$

and set $\chi_n := \exists x \psi_{n-1}(x) \wedge \neg \exists x \psi_n(x)$.

(b) For each $n \geq 1$ there is an FO^3 -formula $\varphi_n(x, y)$ that in digraphs expresses that there is a path of length at most n from x to y . In fact, let

$$\begin{aligned} \varphi_1(x, y) &:= Exy, \\ \varphi_{n+1}(x, y) &:= \varphi_n(x, y) \vee \exists z (Ezy \wedge \exists y (y = z \wedge \varphi_n(x, y))). \end{aligned}$$

Concerning the quantifier rank we can do better than in ψ_n and φ_n : Let φ be an $L_{\infty\omega}^s$ -formula and π a permutation of $1, \dots, s$. By simultaneously replacing both the free and the bound occurrences of v_1, \dots, v_s by $v_{\pi(1)}, \dots, v_{\pi(s)}$ one obtains a formula

$$\varphi \left(\begin{array}{ccc} v_{\pi(1)} & \cdots & v_{\pi(s)} \\ v_1 & \cdots & v_s \end{array} \right).$$

Clearly,

$$\mathcal{A} \models \varphi \left(\begin{array}{ccc} v_{\pi(1)} & \cdots & v_{\pi(s)} \\ v_1 & \cdots & v_s \end{array} \right) [\bar{a}] \quad \text{iff} \quad \mathcal{A} \models \varphi [a_{\pi(1)}, \dots, a_{\pi(s)}].$$

Now in the preceding examples we can replace ψ_n and φ_n by the formulas ψ'_n and φ'_n of quantifier rank $\leq n + 1$, where

$$\begin{aligned} \psi'_0 &:= \psi_0, & \psi'_{n+1} &:= \forall y(y < x \leftrightarrow \bigvee_{i \leq n} \psi'_i \left(\begin{array}{c} yx \\ xy \end{array} \right)), \\ \varphi'_1 &:= \varphi_1, & \varphi'_{n+1} &:= \varphi'_n \vee \exists z(Ezy \wedge \varphi'_n \left(\begin{array}{c} xzy \\ xyz \end{array} \right)). \end{aligned}$$

□

We write $\mathcal{A} \equiv^s \mathcal{B}$, $\mathcal{A} \equiv^{L^s_{\infty\omega}} \mathcal{B}$, and $\mathcal{A} \equiv^s_m \mathcal{B}$ to express that \mathcal{A} and \mathcal{B} satisfy the same sentences in FO^s , in $L^s_{\infty\omega}$, and in FO^s of quantifier rank $\leq m$, respectively.

In view of 3.2.2(b) together with the last remark of 3.2.2 we have

Proposition 3.3.2 *Assume \mathcal{A} and \mathcal{B} are finite. For every $L^s_{\infty\omega}$ -formula φ there is an FO^s -formula ψ with $\text{free}(\psi) \subseteq \text{free}(\varphi)$ such that*

$$\mathcal{A} \models \forall x_1 \dots \forall x_s(\varphi \leftrightarrow \psi) \quad \text{and} \quad \mathcal{B} \models \forall x_1 \dots \forall x_s(\varphi \leftrightarrow \psi).$$

Corollary 3.3.3 *If \mathcal{A} and \mathcal{B} are finite then*

$$\mathcal{A} \equiv^s \mathcal{B} \quad \text{implies} \quad \mathcal{A} \equiv^{L^s_{\infty\omega}} \mathcal{B}.$$

Proof. Given an $L^s_{\infty\omega}$ -sentence φ , choose ψ according to the preceding proposition. Then

$$\begin{aligned} \mathcal{A} \models \varphi &\quad \text{iff} \quad \mathcal{A} \models \psi \\ &\quad \text{iff} \quad \mathcal{B} \models \psi \\ &\quad \text{iff} \quad \mathcal{B} \models \varphi. \end{aligned}$$

□

Clearly, if $\varphi \in \text{FO}^s$ (or, $\varphi \in L^s_{\infty\omega}$) then every subformula of φ contains at most s free variables (namely, at most v_1, \dots, v_s). This property characterizes the formulas of FO^s up to logical equivalence:

Proposition 3.3.4 *Assume $s \geq 1$. If every subformula of $\varphi(v_1, \dots, v_s) \in \text{FO}$ has at most s free variables, then φ is logically equivalent to a formula of FO^s . The statement remains true, if we replace FO and FO^s by $L_{\infty\omega}$ and $L^s_{\infty\omega}$, respectively.*

Proof. By induction (on the quantifier rank) we associate with every formula $\varphi(v_1, \dots, v_s)$ all of whose subformulas have at most s free variables, a formula φ^* with

$$\models \varphi \leftrightarrow \varphi^*, \quad \text{free}(\varphi) = \text{free}(\varphi^*), \quad \varphi^* \in \text{FO}^s.$$

For atomic φ set $\varphi^* := \varphi$, for $\varphi = \neg\chi$ and $\varphi = (\chi_1 \vee \chi_2)$ set $\varphi^* := \neg\chi^*$ and $\varphi^* = (\chi_1^* \vee \chi_2^*)$, respectively. Now, let $\varphi = \exists y\chi$. Then,

$$(*) \quad \text{free}(\chi) \subseteq \{v_1, \dots, v_s, y\} \quad \text{and} \quad \|\text{free}(\chi)\| \leq s.$$

If $y \notin \text{free}(\chi)$ then $\chi = \chi(v_1, \dots, v_s)$ and χ^* is defined by induction hypothesis. We set $\varphi^* := \chi^*$. If $y \in \text{free}(\chi)$ and $y \in \{v_1, \dots, v_s\}$ then, again, $\chi = \chi(v_1, \dots, v_s)$ and χ^* is defined, and we set $\varphi^* := \exists y \chi^*$. Finally suppose $y \in \text{free}(\chi)$ and $y \notin \{v_1, \dots, v_s\}$. Then, by (*), there is an i such that $v_i \notin \text{free}(\chi)$. Set $\chi_0 := \chi \left(\begin{array}{c} v_i \quad y \\ y \quad v_i \end{array} \right)$ (as in the preceding example, $\chi_0 := \chi \left(\begin{array}{c} v_i \quad y \\ y \quad v_i \end{array} \right)$ is obtained from χ by simultaneously replacing *all* occurrences of y and v_i by v_i and y , respectively). Then, $\models \varphi \leftrightarrow \exists v_i \chi_0$, $\text{free}(\chi_0) \subseteq \{v_1, \dots, v_s\}$, and every subformula of χ_0 has at most s free variables. Thus, χ_0^* is defined and we can set $\varphi^* := \exists v_i \chi_0^*$.

3.3.1 Pebble Games

Our next aim is an Ehrenfeucht-Fraïssé type characterization of the equivalence of structures in FO^s and in $\text{L}_{\infty\omega}^s$. To motivate the intended games we look at $\varphi = \exists x \exists y (x < y \wedge \exists x y < x)$ and the orderings $\mathcal{A} := (\{a, b\}, <)$ and $\mathcal{B} := (\{c, d, e\}, <)$ where $a < b$ and $c < d < e$. Since $\mathcal{A} \models \neg \varphi$ and $\mathcal{B} \models \varphi$, the spoiler has a winning strategy in $\text{G}_3(\mathcal{A}, \mathcal{B})$. How is the fact that φ only contains two variables reflected in the course of a play? A play won by the spoiler is given in the following table, where his selections have been underlined.

	\mathcal{A}	\mathcal{B}
first move	a	<u>c</u>
second move	b	<u>d</u>
third move	?	<u>e</u>

There is no third move of the duplicator leading to a partial isomorphism. Apparently, the strategy of the spoiler consists in choosing, for the first two quantifiers $\exists x \exists y$, the elements c for x and d for y in \mathcal{B} in order to have:

$$\mathcal{B} \models (x < y \wedge \exists x y < x)[c, d].$$

The only selections for the duplicator leading to a partial isomorphism are a for x and b for y . Now, for the second quantifier $\exists x$, the spoiler selects in \mathcal{B} the element e , thereby getting a witness for $\mathcal{B} \models \exists x y < x[d]$. Obviously the old value c for x is no longer relevant. Therefore, the play above may be represented more informatively by

	first move		second move		third move	
	\mathcal{A}	\mathcal{B}	\mathcal{A}	\mathcal{B}	\mathcal{A}	\mathcal{B}
x -box	a	<u>c</u>	a	<u>c</u>	?	<u>e</u>
y -box	*	*	b	<u>d</u>	b	<u>d</u>

where the x -boxes and the y -boxes always contain the actual value for x and y , respectively, and $*$ stands for an empty box.

This example motivates how to adapt the Ehrenfeucht-Fraïssé method to our situation: As always, fix a vocabulary τ . By convention, $*$ will not belong to the universe of any structure. For $\bar{a} \in (A \cup \{*\})^s$, $\bar{a} = a_1 \dots a_s$, let $\text{supp}(\bar{a}) := \{i \mid a_i \in A\}$ be the *support* of \bar{a} , and if $a \in A$, let \bar{a}_i^a denote $a_1 \dots a_{i-1} a a_{i+1} \dots a_s$. For $\bar{a} \in (A \cup \{*\})^s$ and $\bar{b} \in (B \cup \{*\})^s$ we say that $\bar{a} \mapsto \bar{b}$ is an *s-partial isomorphism* from \mathcal{A} to \mathcal{B} , if $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$ and $\bar{a}' \mapsto \bar{b}'$ is a partial isomorphism from \mathcal{A} to \mathcal{B} , where \bar{a}' and \bar{b}' are the subsequences of \bar{a} and \bar{b} with indices in the support.

Let \mathcal{A} and \mathcal{B} be structures, $\bar{a} \in (A \cup \{*\})^s$, $\bar{b} \in (B \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$. In the example above we had available a box for each relevant variable. It has become customary to replace the boxes by *pebbles*. Thus, in the *pebble game* $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ we have s pebbles $\alpha_1, \dots, \alpha_s$ for \mathcal{A} and s pebbles β_1, \dots, β_s for \mathcal{B} . Initially, α_i is placed on a_i if $a_i \in A$, and off the board if $a_i = *$, and similarly, β_i is placed on $b_i \in B$ or off the board. Each play consists of m moves. In his j -th move, the spoiler selects a structure, \mathcal{A} or \mathcal{B} , and a pebble for this structure (being off the board or already placed on an element). If he selects \mathcal{A} and α_i , he places α_i on some element of \mathcal{A} , and then the duplicator places β_i on some element of \mathcal{B} . If the spoiler selects \mathcal{B} and β_i , he places β_i on an element of \mathcal{B} and the duplicator places α_i on some element of \mathcal{A} . (Note that there may be several pebbles on the same element.)

The duplicator wins the game if for *each* $j \leq m$ we have that $\bar{e} \mapsto \bar{f}$ is an s -partial isomorphism, where $\bar{e} = e_1 \dots e_s$ are the elements marked by $\alpha_1, \dots, \alpha_s$ after the j -th move ($e_i = *$ in case α_i is off the board) and where $\bar{f} = f_1 \dots f_s$ are the corresponding values given by β_1, \dots, β_s . For $j = 0$ this means that $\bar{a} \mapsto \bar{b}$ is an s -partial isomorphism.

The pebble game $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ with infinitely many moves is defined similarly. We use $G_m^s(\mathcal{A}, \mathcal{B})$ as abbreviation for $G_m^s(\mathcal{A}, * \dots *, \mathcal{B}, * \dots *)$ and $G_\infty^s(\mathcal{A}, \mathcal{B})$ for $G_\infty^s(\mathcal{A}, * \dots *, \mathcal{B}, * \dots *)$.

The following theorem shows that the logics and the games fit together. Here and later, when writing $\mathcal{A} \models \varphi[\bar{a}]$ for $\bar{a} \in (A \cup \{*\})^s$ we tacitly assume that the free variables of φ have indices in $\text{supp}(\bar{a})$ (that is, $i \in \text{supp}(\bar{a})$ whenever $v_i \in \text{free}(\varphi)$).

Theorem 3.3.5 *For structures \mathcal{A} and \mathcal{B} , and for $\bar{a} \in (A \cup \{*\})^s$ and $\bar{b} \in (B \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$ the following hold:*

- (a) \bar{a} satisfies in \mathcal{A} the same FO^s -formulas of quantifier rank $\leq m$ as \bar{b} in \mathcal{B} iff the duplicator wins $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.
- (b) \bar{a} satisfies in \mathcal{A} the same $\text{L}_{\infty\omega}^s$ -formulas as \bar{b} in \mathcal{B} iff the duplicator wins $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.

In particular,

- (c) $\mathcal{A} \equiv_m^s \mathcal{B}$ iff the duplicator wins $G_m^s(\mathcal{A}, \mathcal{B})$.
- (d) $\mathcal{A} \equiv_{\infty\omega}^s \mathcal{B}$ iff the duplicator wins $G_\infty^s(\mathcal{A}, \mathcal{B})$.

Before proving the theorem we give some examples, whose claims are easily verified.

Examples 3.3.6 (a) Let $\tau = \emptyset$ and let \mathcal{A} and \mathcal{B} be τ -structures (i.e., sets). If $\|\mathcal{A}\|, \|\mathcal{B}\| \geq s$ then the duplicator wins $\text{G}_{\infty}^s(\mathcal{A}, \mathcal{B})$ (or, equivalently, $\mathcal{A} \equiv^{\text{L}_{\infty\omega}^s} \mathcal{B}$). Moreover, for arbitrary \mathcal{A} and \mathcal{B} , the duplicator wins $\text{G}_{\infty}^s(\mathcal{A}, \mathcal{B})$ iff he wins $\text{G}_s^s(\mathcal{A}, \mathcal{B})$.

(b) For $l \geq 3$, let \mathcal{G}_l and $\mathcal{G}_l \dot{\cup} \mathcal{G}_l$ be the graphs consisting of one and two cycles, respectively, of length $l + 1$ (cf. 2.3.8). Then the duplicator wins $\text{G}_{\infty}^2(\mathcal{G}_l, \mathcal{G}_l \dot{\cup} \mathcal{G}_l)$ and hence by the theorem, $\mathcal{G}_l \equiv^{\text{L}_{\infty\omega}^2} \mathcal{G}_l \dot{\cup} \mathcal{G}_l$. Note that $\mathcal{G}_l \not\equiv^{\text{L}_{\infty\omega}^3} \mathcal{G}_l \dot{\cup} \mathcal{G}_l$, since – in the notations of Example 3.3.1(b) – the sentence $\forall x \forall y (x = y \vee \bigvee_{n>0} \varphi_n(x, y))$ of $\text{L}_{\infty\omega}^3$ expresses connectivity. In fact, the spoiler wins $\text{G}_{\infty}^3(\mathcal{G}_l, \mathcal{G}_l \dot{\cup} \mathcal{G}_l)$. \square

Further examples are contained in the following exercises.

Exercise 3.3.7 (a) Assume that the relation symbols in $\tau = \{<, \dots\}$ are at most binary. Show that for finite ordered structures \mathcal{A} and \mathcal{B} ,

$$\mathcal{A} \cong \mathcal{B} \quad \text{iff} \quad \text{the duplicator wins } \text{G}_{\infty}^2(\mathcal{A}, \mathcal{B}).$$

(b) For $m \geq s$ show that the duplicator wins $\text{G}_m^s(\mathcal{A}, \mathcal{B})$ iff he wins the game $\text{G}_m^s(\mathcal{A}, \mathcal{B})$ with the additional condition that in the first s moves distinct pebbles have to be chosen (so that after s moves all pebbles are placed on elements).

(c) Suppose that τ is relational and all its relation symbols are of arity $\leq s$. Assume that the duplicator wins $\text{G}_{\infty}^s(\mathcal{A}, \mathcal{B})$ and that $\|\mathcal{A}\| = \|\mathcal{B}\| \leq s + 1$. Show that $\mathcal{A} \cong \mathcal{B}$. \square

A proof of Theorem 3.3.5 can easily be obtained from the corresponding proofs for FO and $\text{L}_{\infty\omega}$: We introduce for FO^s and $\text{L}_{\infty\omega}^s$ the further notions related to the Ehrenfeucht-Fraïssé method (namely, the back and forth properties, the m -isomorphism types, and the set of winning positions), and give the corresponding statements 3.3.9 and 3.3.10 which imply Theorem 3.3.5.

Definition 3.3.8 Structures \mathcal{A} and \mathcal{B} are s - m -isomorphic, $\mathcal{A} \cong_m^s \mathcal{B}$, iff there is a sequence $(I_j)_{j \leq m}$ of nonempty sets of s -partial isomorphisms with the following properties:

(s -forth property) For $j < m$, $\bar{a} \mapsto \bar{b} \in I_{j+1}$, $1 \leq i \leq s$, and $a \in A$ there is $b \in B$ such that $\bar{a} \frac{a}{i} \mapsto \bar{b} \frac{b}{i} \in I_j$.

(s -back property) For $j < m$, $\bar{a} \mapsto \bar{b} \in I_{j+1}$, $1 \leq i \leq s$, and $b \in B$ there is $a \in A$ such that $\bar{a} \frac{a}{i} \mapsto \bar{b} \frac{b}{i} \in I_j$.

We then write $(I_j)_{j \leq m} : \mathcal{A} \cong_m^s \mathcal{B}$.

The notions s -partially isomorphic, $\mathcal{A} \cong_{\text{part}}^s \mathcal{B}$, and $I : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$ are defined similarly. \square

For $m \in \mathbb{N}$, any structure \mathcal{A} , and $\bar{a} \in (A \cup \{*\})^s$, the s - m -isomorphism type $\psi_{\bar{a}}^m (= {}^s\psi_{\mathcal{A}, \bar{a}}^m)$ of \bar{a} in \mathcal{A} is given by:

$$\psi_{\bar{a}}^0(\bar{v}) := \bigwedge \{ \psi \mid \psi \text{ atomic or negated atomic, and } \mathcal{A} \models \psi[\bar{a}] \}$$

(recall that when writing $\mathcal{A} \models \psi[\bar{a}]$ we assume that the free variables of ψ have indices in $\text{supp}(\bar{a})$),

$$\psi_{\bar{a}}^{m+1}(\bar{v}) := \psi_{\bar{a}}^0 \wedge \bigwedge_{1 \leq i \leq s} \bigwedge_{a \in A} \left(\bigwedge \exists v_i \psi_{\bar{a} \frac{a}{i}}^m \wedge \forall v_i \bigvee_{a \in A} \psi_{\bar{a} \frac{a}{i}}^m \right).$$

In particular, $\psi_{\mathcal{A}}^m := \psi_{* \dots *}^m$ is an FO^s -sentence of quantifier rank m .

Finally we introduce the sets $W_m^s(\mathcal{A}, \mathcal{B})$ and $W_\infty^s(\mathcal{A}, \mathcal{B})$ of s -partial isomorphisms corresponding to winning positions in the respective games,

$$\begin{aligned} W_m^s(\mathcal{A}, \mathcal{B}) &:= \{ \bar{a} \mapsto \bar{b} \mid \text{the duplicator wins } G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}) \}, \\ W_\infty^s(\mathcal{A}, \mathcal{B}) &:= \{ \bar{a} \mapsto \bar{b} \mid \text{the duplicator wins } G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}) \}. \end{aligned}$$

Now the following two theorems can be proven completely parallel to those for FO and $L_{\infty\omega}$.

Theorem 3.3.9 *Let $\bar{a} \in (A \cup \{*\})^s$ and $\bar{b} \in (B \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$.*

(a) *The following are equivalent:*

- (i) *The duplicator wins $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.*
- (ii) *$\bar{a} \mapsto \bar{b} \in W_m^s(\mathcal{A}, \mathcal{B})$ and $(W_j^s(\mathcal{A}, \mathcal{B}))_{j \leq m} : \mathcal{A} \cong_m^s \mathcal{B}$.*
- (iii) *There is $(I_j)_{j \leq m}$ with $\bar{a} \mapsto \bar{b} \in I_m$ such that $(I_j)_{j \leq m} : \mathcal{A} \cong_m^s \mathcal{B}$.*
- (iv) *$\mathcal{B} \models \psi_{\bar{a}}^m[\bar{b}]$.*
- (v) *\bar{a} satisfies in \mathcal{A} the same formulas of FO^s of quantifier rank $\leq m$ as \bar{b} in \mathcal{B} .*

(b) *The following are equivalent:*

- (i) *The duplicator wins $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.*
- (ii) *$\bar{a} \mapsto \bar{b} \in W_\infty^s(\mathcal{A}, \mathcal{B})$ and $W_\infty^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$.*
- (iii) *There is I with $\bar{a} \mapsto \bar{b} \in I$ such that $I : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$.*
- (iv) *\bar{a} satisfies in \mathcal{A} the same formulas of $L_{\infty\omega}^s$ as \bar{b} in \mathcal{B} . □*

Corollary 3.3.10 (a) *The following are equivalent:*

- (i) *The duplicator wins $G_m^s(\mathcal{A}, \mathcal{B})$.*
- (ii) *$(W_j^s(\mathcal{A}, \mathcal{B}))_{j \leq m} : \mathcal{A} \cong_m^s \mathcal{B}$.*
- (iii) *$\mathcal{A} \cong_m^s \mathcal{B}$.*
- (iv) *$\mathcal{B} \models \psi_{\mathcal{A}}^m$.*
- (v) *$\mathcal{A} \cong_m^s \mathcal{B}$.*

(b) *The following are equivalent:*

- (i) *The duplicator wins $G_\infty^s(\mathcal{A}, \mathcal{B})$.*
- (ii) *$W_\infty^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$.*
- (iii) *$\mathcal{A} \cong_{\text{part}}^s \mathcal{B}$.*
- (iv) *$\mathcal{A} \cong_{L_{\infty\omega}^s} \mathcal{B}$. □*

Example 3.3.11 For $s \geq 1$ let ϵ_s be the conjunction of the finitely many r -extension axioms with $r \leq s$ (cf. 3.2.11). Clearly, $\epsilon_s \in \text{FO}^s$. We have:

- Every model of ϵ_s has at least s elements.
- Every two models \mathcal{A} and \mathcal{B} of ϵ_s are s -partially isomorphic.

(The last statement is proved as the corresponding one in 3.2.11.) Therefore, by the corollary, any two models of ϵ_s are $L_{\infty\omega}^s$ -equivalent, and we get:

- For every $L_{\infty\omega}^s$ -sentence φ , either $\epsilon_s \models \varphi$ or $\epsilon_s \models \neg\varphi$. □

The next result provides us with an algebraic tool to show that a class of finite structures is not axiomatizable in $L_{\infty\omega}^s$ or in $L_{\infty\omega}^\omega (= \bigcup_{s \geq 1} L_{\infty\omega}^s)$.

Theorem 3.3.12 *Let K be a class of finite structures.*

(a) *For $s \geq 1$ the following are equivalent:*

- (i) *K is not axiomatizable in $L_{\infty\omega}^s$.*
- (ii) *There are finite structures \mathcal{A} and \mathcal{B} such that*

$$\mathcal{A} \in K, \mathcal{B} \notin K, \text{ and } \mathcal{A} \cong_{\text{part}}^s \mathcal{B}.$$

(b) *The following are equivalent:*

- (i) *K is not axiomatizable in $L_{\infty\omega}^\omega$.*
- (ii) *For every $s \geq 1$ there are finite structures \mathcal{A} and \mathcal{B} such that*

$$\mathcal{A} \in K, \mathcal{B} \notin K, \text{ and } \mathcal{A} \cong_{\text{part}}^s \mathcal{B}.$$

Proof. Clearly, (b) follows from (a). To show (ii) \Rightarrow (i) in (a), suppose by contradiction, that $K = \text{Mod}(\varphi)$ for some $\varphi \in L_{\infty\omega}^s$. Choose \mathcal{A} and \mathcal{B} as given by (ii). Then $\mathcal{A} \models \varphi$ (since $\mathcal{A} \in K$), $\mathcal{B} \not\models \varphi$ (since $\mathcal{B} \notin K$), and $\mathcal{A} \equiv_{\infty\omega}^s \mathcal{B}$ (since $\mathcal{A} \cong_{\text{part}}^s \mathcal{B}$), a contradiction.

Conversely, suppose that the condition in (ii) is not satisfied. Then for all finite \mathcal{A} and \mathcal{B} ,

$$(*) \quad \mathcal{A} \in K \text{ and } \mathcal{A} \equiv_{\infty\omega}^s \mathcal{B} \text{ imply } \mathcal{B} \in K.$$

We show that $K = \text{Mod}(\varphi)$ for the $L_{\infty\omega}^s$ -sentence $\varphi := \bigvee_{\mathcal{A} \in K} \bigwedge_{m \geq 0} \psi_{\mathcal{A}}^m$. Clearly, $K \subseteq \text{Mod}(\varphi)$, since $\mathcal{B} \models \bigwedge_{m \geq 0} \psi_{\mathcal{B}}^m$ holds for any \mathcal{B} . To obtain the inclusion $\text{Mod}(\varphi) \subseteq K$, assume that \mathcal{B} is a finite model of φ . Then, for some $\mathcal{A} \in K$ and all $m \geq 0$, we have $\mathcal{B} \models \psi_{\mathcal{A}}^m$ and hence, $\mathcal{A} \equiv_m^s \mathcal{B}$. Thus, $\mathcal{A} \equiv^s \mathcal{B}$. By 3.3.3, we get $\mathcal{A} \equiv_{\infty\omega}^s \mathcal{B}$, and hence by (*), we obtain $\mathcal{B} \in K$. □

The corresponding nonaxiomatizability result for FO^s is part (a) of Exercise 3.3.14 below.

Example 3.3.13 Let τ be the empty vocabulary. The class $\text{EVEN}[\tau]$ is not $L_{\infty\omega}^\omega$ -axiomatizable, since for $s \geq 1$ and structures \mathcal{A} and \mathcal{B} with $\|\mathcal{A}\| = s$ and $\|\mathcal{B}\| = s + 1$, we have $\mathcal{A} \in \text{EVEN}[\tau]$ iff $\mathcal{B} \notin \text{EVEN}[\tau]$; but $\mathcal{A} \equiv_{\infty\omega}^s \mathcal{B}$ by 3.3.6(a). Extend the result to arbitrary τ . □

Exercises 3.3.14 Let K be a class of finite structures. Show:

- (a) The following statements are equivalent:
 (i) K is not axiomatizable in FO^s .
 (ii) For $m \geq 1$ there are finite structures \mathcal{A} and \mathcal{B} such that

$$\mathcal{A} \in K, \mathcal{B} \notin K, \text{ and } \mathcal{A} \cong_m^s \mathcal{B}.$$

- (b) For any global n -ary relation Γ on K with $n \leq s$ the following are equivalent:

- (i) Γ is $L_{\infty\omega}^s$ -definable, i.e. there is an $L_{\infty\omega}^s$ -formula φ such that for $\mathcal{A} \in K$ and $\bar{a} \in A^n$,

$$\mathcal{A} \models \varphi[\bar{a}] \quad \text{iff} \quad \bar{a} \in \Gamma(\mathcal{A}).$$

- (ii) Γ is “closed under the game G_∞^s ”, that means, if $\mathcal{A}, \mathcal{B} \in K$, $\bar{a} \in \Gamma(\mathcal{A})$, $\bar{b} \in B^n$ and the duplicator wins $G_\infty^s(\mathcal{A}, \bar{a} * \dots * \mathcal{B}, \bar{b} * \dots *)$, then $\bar{b} \in \Gamma(\mathcal{B})$. □

Exercise 3.3.15 Assume that τ is relational and contains only unary relation symbols. Fix $s \geq 1$ and let \mathcal{A} and \mathcal{B} be τ structures. Show that $W_s^s(\mathcal{A}, \mathcal{B}) = W_\infty^s(\mathcal{A}, \mathcal{B})$. □

Exercise 3.3.16 For an FO^s -formula $\varphi(\bar{x})$ of quantifier rank $\leq m$ show that

$$\models \varphi \leftrightarrow \bigvee \{ \psi_{\mathcal{A}, \bar{a}}^m \mid \mathcal{A} \text{ a structure, } \bar{a} \in A, \text{ and } \mathcal{A} \models \varphi[\bar{a}] \}$$

(argue as in the proof of 2.2.11).

3.3.2 The s -Invariant of a Structure

In this section we associate with every $s \geq 1$ and every structure \mathcal{A} a structure \mathcal{A}/s , whose isomorphism type captures the $L_{\infty\omega}^s$ -type of \mathcal{A} in a one-to-one fashion (see 3.3.17 below for the precise statement). Later we shall see that there is a uniform way to define an ordering on \mathcal{A}/s . This will enable us to translate some problems from arbitrary structures to ordered ones. For simplicity, let τ be relational (we encourage the reader to also treat vocabularies with constants).

Let \mathcal{A} be a τ -structure. The binary relation \sim defined on A^s by

$$(*) \quad \bar{a} \sim \bar{b} \quad \text{iff} \quad \bar{a} \text{ and } \bar{b} \text{ satisfy the same } L_{\infty\omega}^s\text{-formulas in } \mathcal{A}$$

is an equivalence relation on A^s . By 3.3.9(b), we have $\bar{a} \sim \bar{b}$ iff the duplicator wins $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{A}, \bar{b})$. Let $[\bar{a}]$ denote the equivalence class of \bar{a} and

$$A/s := \{ [\bar{a}] \mid \bar{a} \in A^s \}$$

the set of equivalence classes. We endow A/s with a τ/s -structure \mathcal{A}/s : for every $[\bar{a}] \in A/s$, the relations on A/s capture the properties of \bar{a} in any game $G_\infty^s(\mathcal{A}, \bar{a}, \dots)$. The relation symbols in τ/s (together with their meaning in \mathcal{A}/s) are:

- for every k -ary $R \in \tau \cup \{=\}$ and any i_1, \dots, i_k with $1 \leq i_1, \dots, i_k \leq s$ a unary relation symbol $R_{i_1 \dots i_k}$;

$$R_{i_1 \dots i_k}^{A/s} := \{ [\bar{a}] \mid \bar{a} \in A^s, R^A a_{i_1} \dots a_{i_k} \}$$

(the $R_{i_1 \dots i_k}^{A/s}$ capture the isomorphism type of \bar{a});

- for $i = 1, \dots, s$ a binary relation symbol S_i ;

$$S_i^{A/s} := \{ ([\bar{a}], [\bar{a}']) \mid \bar{a}, \bar{a}' \in A^s, \text{ there is } a \in A \text{ such that } [\bar{a}] = [\bar{a}'/i] \}$$

($S_i^{A/s}$ encodes the possible moves of the i -th pebble).

A/s is called the s -invariant of \mathcal{A} . It captures the $L_{\infty\omega}^s$ -theory of \mathcal{A} :

Theorem 3.3.17 *For structures \mathcal{A} and \mathcal{B} ,*

$$\mathcal{A} \equiv_{L_{\infty\omega}^s} \mathcal{B} \quad \text{iff} \quad \mathcal{A}/s \cong \mathcal{B}/s.$$

Proof. Suppose first that \mathcal{A}/s and \mathcal{B}/s are isomorphic and $\pi : \mathcal{A}/s \cong \mathcal{B}/s$. Set

$$I := \{ \bar{a} \mapsto \bar{b} \mid \bar{a} \in A^s, \bar{b} \in B^s, \pi([\bar{a}]) = [\bar{b}] \}.$$

We show that $I : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$ and hence, $\mathcal{A} \equiv_{L_{\infty\omega}^s} \mathcal{B}$. I is a nonempty set of s -partial isomorphisms (use the relations $R_{i_1 \dots i_k}$). To show, say, that I has the s -forth property, assume that $\bar{a} \mapsto \bar{b} \in I$, $1 \leq i \leq s$, and $a \in A$. Then $S_i^{A/s}[\bar{a}][\bar{a}'/i]$, hence $S_i^{B/s}[\bar{b}]\pi([\bar{a}'/i])$. By the definition of $S_i^{B/s}$, there is $b \in B$ such that $[\bar{b}'/i] = \pi([\bar{a}'/i])$, hence $\bar{a}'/i \mapsto \bar{b}'/i \in I$.

Conversely suppose that $\mathcal{A} \equiv_{L_{\infty\omega}^s} \mathcal{B}$. Then $W_{\infty}^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$. For $\bar{a} \in A^s$ and $\bar{b} \in B^s$ set

$$\pi([\bar{a}]) := [\bar{b}] \quad \text{iff} \quad \bar{a} \mapsto \bar{b} \in W_{\infty}^s(\mathcal{A}, \mathcal{B}).$$

Hence, by 3.3.9(b),

$$\pi([\bar{a}]) = [\bar{b}] \quad \text{iff} \quad \bar{a} \text{ in } \mathcal{A} \text{ satisfies the same } L_{\infty\omega}^s\text{-formulas as } \bar{b} \text{ in } \mathcal{B}.$$

By this equivalence and by the definition (*) of the equivalence relation, π is well-defined and injective; moreover, $\text{do}(\pi) = \mathcal{A}/s$ by the s -forth property of $W_{\infty}^s(\mathcal{A}, \mathcal{B})$ and $\text{rg}(\pi) = \mathcal{B}/s$ by the s -back property. Obviously, π is compatible with the interpretations of the $R_{i_1 \dots i_k}$, and also with the interpretations of the S_i (use once more the s -back and s -forth properties of $W_{\infty}^s(\mathcal{A}, \mathcal{B})$). Therefore, $\pi : \mathcal{A}/s \cong \mathcal{B}/s$. \square

We close this section by showing that for finite structures we can replace $L_{\infty\omega}^s$ by FO^s . In fact, in view of 3.3.2 we have

Proposition 3.3.18 *Let \mathcal{A} and \mathcal{B} be finite structures, $\bar{a} \in (A \cup \{*\})^s$ and $\bar{b} \in (B \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$. If*

$$\text{for all } \varphi \in \text{FO}^s \text{ } (\mathcal{A} \models \varphi[\bar{a}] \text{ iff } \mathcal{B} \models \varphi[\bar{b}])$$

then

$$\text{for all } \varphi \in \text{L}_{\infty\omega}^s \text{ } (\mathcal{A} \models \varphi[\bar{a}] \text{ iff } \mathcal{B} \models \varphi[\bar{b}]). \quad \square$$

Corollary 3.3.19 (a) *Let \mathcal{A} be a finite structure and \sim be defined as in (*) above. Then, for $\bar{a}, \bar{b} \in A^s$,*

$$\bar{a} \sim \bar{b} \text{ iff } \bar{a} \text{ and } \bar{b} \text{ satisfy the same } \text{FO}^s\text{-sentences.}$$

(b) *For finite structures \mathcal{A} and \mathcal{B} ,*

$$\mathcal{A} \equiv^s \mathcal{B} \text{ iff } \mathcal{A}/s \cong \mathcal{B}/s.$$

Proof. (a) is immediate from the preceding proposition and (b) follows from 3.3.17 by 3.3.3. \square

3.3.3 Scott Formulas

For a finite structure \mathcal{A} we saw in 2.2.10 that a single FO-sentence, e.g. $\varphi_{\mathcal{A}}^{||\mathcal{A}||+1}$, characterizes \mathcal{A} up to (isomorphisms and therefore up to) $\text{L}_{\infty\omega}$ -equivalence. We are going to show a similar result for FO^s and $\text{L}_{\infty\omega}^s$, in this way strengthening 3.3.3.

Proposition 3.3.20 *Let \mathcal{A} and \mathcal{B} be structures. Then*

- (a) $W_0^s(\mathcal{A}, \mathcal{B}) \supseteq W_1^s(\mathcal{A}, \mathcal{B}) \supseteq \dots$
- (b) *If \mathcal{A} and \mathcal{B} are finite then there is an $m \leq (||\mathcal{A}|| + 1)^s \cdot (||\mathcal{B}|| + 1)^s$ such that $W_m^s(\mathcal{A}, \mathcal{B}) = W_{m+1}^s(\mathcal{A}, \mathcal{B})$.*
- (c) *For $m \geq 0$, if $W_m^s(\mathcal{A}, \mathcal{B}) = W_{m+1}^s(\mathcal{A}, \mathcal{B})$ and $W_m^s(\mathcal{A}, \mathcal{B})$ is nonempty, then $W_m^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}} \mathcal{B}$.*

Proof. (a) follows immediately from the definition of the $W_j^s(\mathcal{A}, \mathcal{B})$. (b) follows from (a), since there are at most $(||\mathcal{A}|| + 1)^s \cdot (||\mathcal{B}|| + 1)^s$ s -partial isomorphisms from \mathcal{A} to \mathcal{B} .

(c) Suppose that $W_m^s(\mathcal{A}, \mathcal{B}) = W_{m+1}^s(\mathcal{A}, \mathcal{B})$. Then $W_m^s(\mathcal{A}, \mathcal{B})$ has the s -back and the s -forth property: To show, say, the s -forth property, let $\bar{a} \mapsto \bar{b} \in W_m^s(\mathcal{A}, \mathcal{B})$, $1 \leq i \leq s$, and $a \in A$. By assumption, $\bar{a} \mapsto \bar{b} \in W_{m+1}^s(\mathcal{A}, \mathcal{B})$, and therefore there is $b \in B$ such that $\bar{a}_i^a \mapsto \bar{b}_i^b \in W_m^s(\mathcal{A}, \mathcal{B})$. If, in addition, $W_m^s(\mathcal{A}, \mathcal{B}) \neq \emptyset$, we altogether have $W_m^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}} \mathcal{B}$. \square

Fix a finite structure \mathcal{A} and let \bar{a}, \bar{b} range over $(A \cup \{*\})^s$. By the proposition we know that

$$W_0^s(\mathcal{A}, \mathcal{A}) \supseteq W_1^s(\mathcal{A}, \mathcal{A}) \supseteq \dots \supseteq W_m^s(\mathcal{A}, \mathcal{A}) \supseteq \dots$$

and that there must be a j with $W_j^s(\mathcal{A}, \mathcal{A}) = W_{j+1}^s(\mathcal{A}, \mathcal{A})$. The minimal j having this property is called the s -rank $r(\mathcal{A})$ of \mathcal{A} , $r(\mathcal{A}) = r(s, \mathcal{A})$. For given \bar{a} , the formula

$$\sigma_{\bar{a}} := \psi_{\bar{a}}^{r(\mathcal{A})} \wedge \bigwedge_{\bar{b} \in (\mathcal{A} \cup \{*\})^s} \forall v_1 \dots \forall v_s (\psi_{\bar{b}}^{r(\mathcal{A})} \rightarrow \psi_{\bar{b}}^{r(\mathcal{A})+1})$$

(more exactly, $\sigma_{\bar{a}} = {}^s\sigma_{\mathcal{A}, \bar{a}}$) is called the s -Scott formula of \bar{a} in \mathcal{A} . It is an FO^s -formula of quantifier rank $r(\mathcal{A}) + 1 + s$. In particular, $\sigma_{\mathcal{A}} := \sigma_{*\dots*}$ is an FO^s -sentence. It captures the whole $\text{L}_{\infty\omega}^s$ -theory of \mathcal{A} :

Theorem 3.3.21 *Let \mathcal{A} be finite.*

(a) *For any structure \mathcal{B} ,*

$$\mathcal{B} \models \sigma_{\mathcal{A}} \quad \text{iff} \quad \mathcal{A} \equiv_{\text{L}_{\infty\omega}^s} \mathcal{B}.$$

(b) *For $\bar{a} \in (\mathcal{A} \cup \{*\})^s$, any structure \mathcal{B} and $\bar{b} \in (\mathcal{B} \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$,*

$$\mathcal{B} \models \sigma_{\bar{a}}[\bar{b}] \quad \text{iff} \quad \bar{a} \text{ satisfies in } \mathcal{A} \text{ the same } \text{L}_{\infty\omega}^s\text{-formulas as } \bar{b} \text{ in } \mathcal{B}.$$

Proof. We restrict ourselves to (a). Since $\mathcal{A} \models \sigma_{\mathcal{A}}$, we have that $\mathcal{A} \equiv_{\text{L}_{\infty\omega}^s} \mathcal{B}$ implies $\mathcal{B} \models \sigma_{\mathcal{A}}$. Now suppose that $\mathcal{B} \models \sigma_{\mathcal{A}}$, that is,

$$\mathcal{B} \models \psi_{\mathcal{A}}^{r(\mathcal{A})} \wedge \bigwedge_{\bar{b} \in (\mathcal{A} \cup \{*\})^s} \forall v_1 \dots \forall v_s (\psi_{\bar{b}}^{r(\mathcal{A})} \rightarrow \psi_{\bar{b}}^{r(\mathcal{A})+1}).$$

Since $\mathcal{B} \models \psi_{\mathcal{A}}^{r(\mathcal{A})}$, we get $*\dots* \mapsto *\dots* \in W_{r(\mathcal{A})}^s(\mathcal{A}, \mathcal{B})$ (by 3.3.9(a)). By the validity of the second conjunct in \mathcal{B} , we have $W_{r(\mathcal{A})}^s(\mathcal{A}, \mathcal{B}) \subseteq W_{r(\mathcal{A})+1}^s(\mathcal{A}, \mathcal{B})$ and hence, $W_{r(\mathcal{A})}^s(\mathcal{A}, \mathcal{B}) = W_{r(\mathcal{A})+1}^s(\mathcal{A}, \mathcal{B})$. Therefore, by 3.3.20(c), $W_{r(\mathcal{A})}^s(\mathcal{A}, \mathcal{B}) : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}$ and thus, $\mathcal{A} \equiv_{\text{L}_{\infty\omega}^s} \mathcal{B}$. \square

Corollary 3.3.22 *In the finite, each $\text{L}_{\infty\omega}^s$ -formula φ is equivalent to a countable disjunction of FO^s -formulas. In fact, φ is equivalent to the $\text{L}_{\infty\omega}^s$ -formula $\bigvee \{\sigma_{\bar{a}} \mid \mathcal{A} \text{ finite, } \bar{a} \in \mathcal{A}, \mathcal{A} \models \varphi[\bar{a}]\}$. Moreover, if K is any class of finite structures, then φ and $\bigvee \{\sigma_{\bar{a}} \mid \mathcal{A} \in K, \bar{a} \in \mathcal{A}, \mathcal{A} \models \varphi[\bar{a}]\}$ are equivalent in all structures of K . \square*

Exercise 3.3.23 For finite \mathcal{A} and \mathcal{B} show that $W_j^s(\mathcal{A}, \mathcal{A}) = W_{j+1}^s(\mathcal{A}, \mathcal{A})$ and $\mathcal{A} \equiv_{\text{L}_{\infty\omega}^s} \mathcal{B}$ imply $W_j^s(\mathcal{B}, \mathcal{B}) = W_{j+1}^s(\mathcal{B}, \mathcal{B})$. Conclude: If $\mathcal{A} \equiv_{\text{L}_{\infty\omega}^s} \mathcal{B}$ then $r(\mathcal{A}) = r(\mathcal{B})$. \square

As an application of the formulas $\sigma_{\bar{a}}$ we give a condition for $\text{L}_{\infty\omega}^s$ and FO^s to coincide in expressive power.

Let K be a class of finite structures. We say that K is s -bounded if the set $\{r(\mathcal{A}) \mid \mathcal{A} \in K\}$ of s -ranks of structures in K is bounded. The class K is *bounded* if it is s -bounded for every $s \geq 1$.

Theorem 3.3.24 *Let K be a class of finite structures.*

- (a) *For $s \geq 1$ the following are equivalent:*
 - (i) *K is s -bounded.*
 - (ii) *On K , every $L_{\infty\omega}^s$ -formula is equivalent to an FO^s -formula.*
 - (iii) *On K , every $L_{\infty\omega}^s$ -formula is equivalent to an FO -formula.*
- (b) *K is bounded iff FO and $L_{\infty\omega}^\omega$ have the same expressive power on K .*

Proof. As (b) is a consequence of (a), it suffices to prove (a). First suppose that K is s -bounded and set $m := \sup\{r(\mathcal{A}) \mid \mathcal{A} \in K\} < \infty$. Thus, for $\mathcal{A} \in K$ and \bar{a} in \mathcal{A} , the quantifier rank of $\sigma_{\bar{a}}$ is $\leq m + s + 1$. Let φ be any $L_{\infty\omega}^s$ -formula. Then the disjunction in the preceding corollary is a disjunction of formulas of quantifier rank $\leq m + s + 1$ and hence, it is a finite one. This shows that (i) implies (ii). The implication from (ii) to (iii) is trivial. To show that (iii) implies (i) assume, by contradiction, that K is not s -bounded. Let $\mathcal{A}_0, \mathcal{A}_1, \dots$ be structures in K of pairwise distinct s -rank. For $M \subseteq \mathbb{N}$ let

$$\varphi_M := \bigvee \{\sigma_{\mathcal{A}_i} \mid i \in M\}.$$

By 3.3.23, if $L, M \subseteq \mathbb{N}$ and $L \neq M$ then not $K \models \varphi_L \leftrightarrow \varphi_M$.⁴ Hence on K , $L_{\infty\omega}^s$ contains uncountably many pairwise nonequivalent sentences and is therefore more expressive than FO . \square

Example 3.3.25 (a) Suppose $\tau = \emptyset$ and let \mathcal{A} be a τ -structure. Then $W_0^s(\mathcal{A}, \mathcal{A}) = W_1^s(\mathcal{A}, \mathcal{A})$. Hence, $r(\mathcal{A}) = 0$ and the class K of finite τ -structures is bounded. Therefore, FO and $L_{\infty\omega}^\omega$ have the same expressive power on K .

(b) By Exercise 3.3.15, the class K of all finite τ -structures is bounded, if τ contains only unary relation symbols. Therefore, FO and $L_{\infty\omega}^\omega$ have the same expressive power on K . \square

Exercise 3.3.26 $L_{\infty\omega}^\omega$ is stronger than FO on the class of finite orderings and on the class of graphs. (Hint: Use 3.3.1.) \square

3.4 Logics with Counting Quantifiers

To express in first-order logic that there are, say, seven elements with the property $\varphi(x)$ we need, in general, at least seven quantifiers:

$$\exists x_1 \dots \exists x_7 (\varphi(x_1) \wedge \dots \wedge \varphi(x_7) \wedge \bigwedge_{1 \leq i < j \leq 7} \neg x_i = x_j)$$

(by 3.3.6(a) we see that in case $\varphi(x) := x = x$ we really need seven quantifiers).

⁴ Recall that $K \models \varphi$ means that every structure in K is a model of φ .

Let $\text{FO}(\mathcal{C})$, so-called *first-order logic with counting quantifiers*, and $\text{L}_{\infty\omega}(\mathcal{C})$ (by short, $\text{C}_{\infty\omega}$), so-called *L_{∞ω} with counting quantifiers*, be the logics obtained from FO and $\text{L}_{\infty\omega}$, respectively, by adding, for every $l \geq 1$, a new quantifier $\exists^{\geq l}$ with the intended interpretation “there are at least l ”. More precisely, extend the calculus of formulas for first-order or infinitary logic by the following rule:

– If φ is a formula and $l \geq 1$ then $\exists^{\geq l}x\varphi$ is a formula

(here, $\exists^{\geq l}x$ is considered as a new quantifier (binding x) and not as an abbreviation as in 1.B). Fix the interpretation of these quantifiers by adding for $\varphi = \varphi(\bar{x}, x)$ and $\bar{a} \in A$ the clause

$$\mathcal{A} \models \exists^{\geq l}x\varphi[\bar{a}] \quad \text{iff} \quad \|\{b \in A \mid \mathcal{A} \models \varphi[\bar{a}, b]\}\| \geq l.$$

Since the quantifiers $\exists^{\geq l}$ are first-order definable, the languages $\text{FO}(\mathcal{C})$ and $\text{C}_{\infty\omega}$ have the same expressive power as FO and $\text{L}_{\infty\omega}$, respectively. The situation changes if we restrict ourselves to $\text{FO}(\mathcal{C})^s$ and $\text{C}_{\infty\omega}^s$, which are the fragments consisting of the formulas with variables among v_1, \dots, v_s . For example, $\exists^{\geq 7}xx = x$ is a sentence in $\text{FO}(\mathcal{C})^1$ not equivalent to any sentence in FO^1 . And $\bigvee_{l \geq 1}(\exists^{\geq 2l}xx = x \wedge \neg \exists^{\geq 2l+1}xx = x)$ is a $\text{C}_{\infty\omega}^1$ -sentence axiomatizing the class $\text{EVEN}[\tau]$ of structures of even cardinality that (by 3.3.13) is not equivalent to any sentence of $\text{L}_{\infty\omega}^{\omega}$. We set $\text{C}_{\infty\omega}^{\omega} := \bigcup_{s \geq 1} \text{C}_{\infty\omega}^s$.

For $l \geq 1$, set $\exists^=l x\varphi := \exists^{\geq l}x\varphi \wedge \neg \exists^{\geq l+1}x\varphi$, and let $\exists^=0 x\varphi := \forall x\neg\varphi$. Then $\exists^{\geq l}x\varphi$ is equivalent to $\neg \bigvee_{j < l} \exists^=j x\varphi$. Hence, we would obtain logics of the same expressive power when adding the quantifiers $\exists^=l$ instead of $\exists^{\geq l}$.

Examples 3.4.1 (a) Suppose \mathcal{A} and \mathcal{B} are finite structures. If $\mathcal{A} \equiv^{\text{FO}(\mathcal{C})^1} \mathcal{B}$, that is, if \mathcal{A} and \mathcal{B} satisfy the same sentences of $\text{FO}(\mathcal{C})^1$, then $\|\mathcal{A}\| = \|\mathcal{B}\|$ (note that $\exists^{\|\mathcal{A}\|}xx = x$ is a sentence in $\text{FO}(\mathcal{C})^1$).

(b) Let $\tau = \{<\}$. The finite models of the sentence of $\text{FO}[\tau]$

$$\begin{aligned} & \forall x\neg x < x \wedge \forall x\forall y\forall z((x < y \wedge y < z) \rightarrow x < z) \wedge \\ & \forall x\forall y\forall z((y < x \wedge z < x) \rightarrow (y < z \vee y = z \vee z < y)) \end{aligned}$$

($<$ is irreflexive and transitive, and the predecessors of any element are linearly ordered) are called finite *<-forests*. For a $<$ -forest \mathcal{A} and $a \in \mathcal{A}$ the height $h_{\mathcal{A}}(a)$ is defined by

$$h_{\mathcal{A}}(a) := \|\{b \in \mathcal{A} \mid b < a\}\|,$$

and the height $h(\mathcal{A})$ by

$$h(\mathcal{A}) := \max\{h_{\mathcal{A}}(a) \mid a \in \mathcal{A}\}.$$

The element a is a *root* if $h_{\mathcal{A}}(a) = 0$. Every finite $<$ -forest can be characterized, up to isomorphism, in $\text{FO}(\mathcal{C})^2$, i.e., for every finite $<$ -forest \mathcal{A} there is a sentence φ in $\text{FO}(\mathcal{C})^2$ such that for all finite $<$ -forests \mathcal{B} ,

$$\mathcal{B} \models \varphi \quad \text{iff} \quad \mathcal{B} \cong \mathcal{A}.$$

To prove this one shows by induction on the height, that for $<$ -forests \mathcal{A} with exactly one root there is a formula $\psi_{\mathcal{A}}(x)$ in $\text{FO}(\mathcal{C})^2$ such that for any $<$ -forest \mathcal{B} and $b \in B$ one has

$$\mathcal{B} \models \psi_{\mathcal{A}}[b] \quad \text{iff} \quad \mathcal{B}_b \cong \mathcal{A},$$

where \mathcal{B}_b is the substructure of \mathcal{B} with universe $\{b' \in B \mid b = b' \vee b < b'\}$. In the induction step, for \mathcal{A} with root a , $\psi_{\mathcal{A}}(x)$ gives the number of elements of A and, for any isomorphism type of some \mathcal{A}_b with $b \in A \setminus \{a\}$, says how many trees \mathcal{A}_c with $c \in A \setminus \{a\}$ are of this type.

(c) For $s \geq 1$ there are $<$ -forests \mathcal{A} and \mathcal{B} that satisfy the same sentences in FO^s but are not isomorphic, e.g., $<$ -forests consisting only of roots, the first one having s roots, the second one $s + 1$ roots. \square

We conclude this section by adapting the Ehrenfeucht-Fraïssé method to the case of counting quantifiers and by demonstrating with an example its value for investigating the expressive power of the logics $\text{FO}(\mathcal{C})^s$ and $\text{C}_{\infty\omega}^s$.

In the corresponding pebble games $\text{C-G}_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ with m moves and $\text{C-G}_{\infty}^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$ with infinitely many moves, each move now consists of two steps:

1. The spoiler chooses one of the two structures, say \mathcal{A} , and a corresponding pebble, say α_i . He then chooses a subset X of A . The duplicator must answer with a subset Y of B of the same cardinality as X .
2. The spoiler places β_i on some element $b \in Y$. The duplicator answers by placing α_i on some $a \in X$ (X and Y can now be forgotten).

The definition for winning is given as in the previous pebble games; it only takes into consideration the chosen elements, not the subsets.

To understand what is going on in the two steps of a move suppose that the spoiler attempts to show that

$$\mathcal{A} \models \exists^{\geq l} x \varphi(x), \quad \text{but not } \mathcal{B} \models \exists^{\geq l} x \varphi(x).$$

He chooses a subset X consisting of l elements witnessing that $\mathcal{A} \models \exists^{\geq l} x \varphi(x)$. The duplicator claims that the elements of the subset Y witness that $\mathcal{B} \models \exists^{\geq l} x \varphi(x)$. According to the spoiler's conviction there is a $b \in Y$ with not $\mathcal{B} \models \varphi[b]$. The duplicator means that some element a in X behaves as b .

Exercise 3.4.2 Let \mathcal{A} and \mathcal{B} be finite $<$ -forests and assume that the duplicator has a winning strategy in $\text{C-G}_{\infty}^2(\mathcal{A}, \mathcal{B})$. Show that \mathcal{A} and \mathcal{B} are isomorphic. \square

The reader will encounter no difficulties when trying to prove the following theorem like corresponding preceding results.

Theorem 3.4.3 *Let \mathcal{A} and \mathcal{B} be structures, $\bar{a} \in (A \cup \{*\})^s$ and $\bar{b} \in (B \cup \{*\})^s$ with $\text{supp}(\bar{a}) = \text{supp}(\bar{b})$.*

(a) *The following are equivalent:*

(i) *For all $\varphi(\bar{x}) \in C_{\infty\omega}^s$: $\mathcal{A} \models \varphi[\bar{a}]$ iff $\mathcal{B} \models \varphi[\bar{b}]$.*

(ii) *The duplicator wins $C\text{-G}_{\infty}^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$.*

(b) *$\mathcal{A} \equiv_{C_{\infty\omega}^s} \mathcal{B}$ iff the duplicator wins $C\text{-G}_{\infty}^s(\mathcal{A}, \mathcal{B})$.* □

We now use this theorem to show that for coloured graphs the expressive power of $C_{\infty\omega}^2$ corresponds to a natural graph-theoretical property.

Let C_1, C_2, \dots be unary relation symbols, the “colour relations”. A *coloured graph* is, for some r , an $\{E, C_1, \dots, C_r\}$ -structure \mathcal{G} , where for $\mathcal{G} = (G, E^G, C_1^G, \dots, C_r^G)$ the following holds:

- (G, E^G) is a graph;
- $C_1^G \dot{\cup} \dots \dot{\cup} C_r^G = G$, that is, each vertex satisfies exactly one colour relation.

For $a \in G$ the *colour type* $\text{ct}(a)$ is defined as

$$\text{ct}(a) := (i, n_1, \dots, n_r),$$

where

$$a \in C_i^G \quad \text{and} \quad n_j := \|\{b \in C_j^G \mid E^G ab\}\|.$$

\mathcal{G} is *stable* if for $a, b \in G$,

$$\text{ct}(a) = \text{ct}(b) \quad \text{iff} \quad a, b \in C_i \text{ for some } i.$$

The proof of the following statement is straightforward.

Proposition 3.4.4 *Let $\mathcal{G} = (G, E^G, C_1^G, \dots, C_r^G)$ be a stable coloured graph and $a, b \in G$. Then the following are equivalent:*

(i) *For $j = 1, \dots, r$: $a \in C_j$ iff $b \in C_j$.*

(ii) *The duplicator has a winning strategy in the game $C\text{-G}_{\infty}^2(\mathcal{G}, a*, \mathcal{G}, b*)$.* □

We introduce a process of colour refinement leading from a coloured graph $\mathcal{G} = (G, E^G, C_1^G, \dots, C_r^G)$ to a stable coloured graph: Let $m := \|\{\text{ct}(a) \mid a \in G\}\|$ and order the set $\{\text{ct}(a) \mid a \in G\}$ lexicographically. Set $\mathcal{G}' := (G, E^G, C'_1, \dots, C'_m)$, where C'_k is the set of elements $a \in G$ such that $\text{ct}(a)$ is the k -th element in this ordering. Clearly,

(*) $\text{each } C'_i \text{ is the union of some } C_k^G.$

Moreover, if C'_k is the set of elements of colour type (i, n_1, \dots, n_r) , then $C'_k = \{a \in G \mid \mathcal{G} \models (C_i x \wedge \bigwedge_{j=1, \dots, r} \exists^{=n_j} y (E x y \wedge C_j y)) [a]\}$, that is

⁵ Of course, the corresponding results for $\text{FO}(C)^s$ hold, too.

- (*) each colour class of \mathcal{G}' is definable in \mathcal{G} by a formula of $C_{\infty\omega}^2$ of quantifier rank ≤ 1 ,

where we extend the definition of quantifier rank for first-order logic by the clause $\text{qr}(\exists^{\geq l} x\varphi) := 1 + \text{qr}(\varphi)$.

Obviously, \mathcal{G} is stable if $m = r$, i.e., if there is no proper colour refinement. If \mathcal{G}' is not stable, define $\mathcal{G}^{(2)} := (\mathcal{G}')'$. By (*), we get that after finitely many steps, say after n steps, we must reach a stable coloured graph $\mathcal{G}^{(n)}$, the *stable coloured refinement* of \mathcal{G} . A simple induction, using (*), shows that each colour class of $\mathcal{G}^{(n)}$ is definable by a $C_{\infty\omega}^2$ -formula of quantifier rank $\leq n$. This fact yields (ii) \Rightarrow (i) of the following theorem.

Theorem 3.4.5 *For elements a and b of a coloured graph \mathcal{G} the following are equivalent:*

- (i) a, b are in the same colour class of the stable coloured refinement of \mathcal{G} .
(ii) For all $\varphi(x) \in C_{\infty\omega}^2$,

$$\mathcal{G} \models \varphi[a] \quad \text{iff} \quad \mathcal{G} \models \varphi[b].$$

To prove (i) \Rightarrow (ii), use the preceding proposition (note that by (*), a winning strategy for the coloured refinement of \mathcal{G} is a winning strategy for \mathcal{G}). \square

For a graph $\mathcal{G} = (G, E^G)$ let the stable coloured refinement be that of the coloured graph (G, E^G, G) . Then

Corollary 3.4.6 *For elements a and b of a graph \mathcal{G} the following are equivalent:*

- (i) a, b are in the same colour class of the stable coloured refinement of \mathcal{G} .
(ii) For all $\varphi(x) \in C_{\infty\omega}^2$,

$$\mathcal{G} \models \varphi[a] \quad \text{iff} \quad \mathcal{G} \models \varphi[b]. \quad \square$$

3.5 Failure of Classical Theorems in the Finite

We have seen that in the finite the compactness theorem for first-order logic fails and we shall see in Chapter 7 that there is no sound and complete proof calculus.

These facts raise the question whether and why first-order logic can serve as a useful means for the investigation of finite structures. The question demands for an answer even more as other central methods of classical model theory (such as the use of ultraproducts or saturated structures) become useless and, as we shall see in this section, further important results fail when restricted to finite structures. The examples we give include Beth's definability theorem, Craig's interpolation theorem, and some preservation

theorems. Nevertheless, we hope to convince the reader that new methods and results intrinsic to the finite compensate for this loss. As a new aspect we mention the stronger impact of combinatorics which, in particular, will become apparent in connection with probabilities (see Chapter 4). Moreover, the restriction to the finite motivates the use of other languages, for example languages that are able to grasp notions of recursion or induction, in this way building a bridge to computational aspects (see Chapter 7).

Recall that \models denotes the consequence relation with respect to arbitrary structures and \models_{fin} that for finite structures; so $\varphi \models_{\text{fin}} \psi$ means that every finite model of φ is a model of ψ .

The Beth Property and the Interpolation Property

Let \mathcal{L} be any logic considered so far, e.g. FO, $L_{\infty\omega}, \dots$. Let R be an n -ary relation symbol not contained in the vocabulary τ . An $\mathcal{L}[\tau \cup \{R\}]$ -sentence φ defines R implicitly (in the finite), if every (finite) τ -structure \mathcal{A} has at most one expansion $(\mathcal{A}, R^{\mathcal{A}})$ to a $\tau \cup \{R\}$ -structure satisfying φ . We say that R is explicitly definable (in the finite) relative to φ , if there is an $\mathcal{L}[\tau]$ -formula $\psi(\bar{x})$ such that

$$\varphi \models_{(\text{fin})} \forall \bar{x}(R\bar{x} \leftrightarrow \psi(\bar{x})).$$

Obviously, if R is explicitly definable relative to φ then φ defines R implicitly. We say that \mathcal{L} has the Beth property (in the finite), if the converse holds, i.e., whenever an \mathcal{L} -sentence φ defines a relation symbol implicitly (in the finite), then there is an explicit definition of it (in the finite) relative to φ .

Beth’s theorem states that first-order logic has the Beth property. It exhibits a certain balance between syntax and semantics: every implicitly definable, i.e., “semantically” definable, relation has an explicit, a “syntactic” definition.

Proposition 3.5.1 *First-order logic does not have the Beth property in the finite.*

Proof. We consider orderings in the vocabulary $\tau := \{<, S, \min, \max\}$. Let R be a unary relation symbol and let φ be the conjunction of the ordering axioms and of the following sentence fixing R as the set of even points,

$$\neg R \min \wedge \forall x \forall y (Sxy \rightarrow (Rx \leftrightarrow \neg Ry)).$$

Clearly, φ defines R implicitly in the finite. Suppose

$$\varphi \models_{\text{fin}} \forall x (Rx \leftrightarrow \psi(x))$$

for some FO $[\tau]$ -formula $\psi(x)$. Then $\psi(\max)$ together with the ordering axioms would define the class of finite orderings of even cardinality, contradicting 2.3.6. □

Closely related to the Beth property is the interpolation property (or, Craig property) (more precisely, the Beth property is a consequence of the interpolation property): The logic \mathcal{L} has the *interpolation property (in the finite)* iff for all vocabularies σ and τ and any \mathcal{L} -sentences φ and ψ in the vocabularies σ and τ , respectively, such that $\varphi \models_{(\text{fin})} \psi$, there is an *interpolant*, that is, an $\mathcal{L}[\sigma \cap \tau]$ -sentence χ such that

$$\varphi \models_{(\text{fin})} \chi \quad \text{and} \quad \chi \models_{(\text{fin})} \psi.$$

Craig's theorem states that first-order logic has the interpolation property. We shall see that in the finite the interpolation property fails for first-order logic. Our counterexample deals with a special case of the interpolation property that will be important later. Given a logic \mathcal{L} , some classes K of finite τ -structures happen to be axiomatizable in \mathcal{L} , if we equip the structures in K with an arbitrary ordering, that is, there is a sentence φ in the vocabulary $\tau \dot{\cup} \{<\}$ such that for

$$K_{<} := \{(\mathcal{A}, <) \mid \mathcal{A} \in K, < \text{ an ordering on } A\},$$

we have

$$K_{<} = \text{Mod}(\varphi).$$

The logic \mathcal{L} is said to be *closed under order-invariant sentences in the finite*, whenever in this situation there is an $\mathcal{L}[\tau]$ -sentence ψ such that $\text{Mod}(\psi) = K$.

If $\varphi = \varphi(<)$ axiomatizes $K_{<}$ and $<'$ is a new binary relation symbol, then

$$\varphi(<) \models_{\text{fin}} ("<' \text{ is an ordering}" \rightarrow \varphi(<')).$$

Clearly, if ψ is any interpolant then $\text{Mod}(\psi) = K$. Hence, a logic with the interpolation property is closed under order-invariant sentences in the finite. In particular, in the following proposition part (b) follows from part (a).

Proposition 3.5.2 (a) *First-order logic is not closed under order-invariant sentences in the finite.*

(b) *First-order logic does not have the interpolation property in the finite.*

We sketch a *proof* of (a). Let K be the class of finite Boolean algebras with an even number of atoms. Using the Ehrenfeucht-Fraïssé method, one can show that K is not axiomatizable in first-order logic. However, $K_{<}$ is axiomatizable in first-order logic. In fact, let φ be the conjunction of the axioms for Boolean algebras and the axioms for orderings and of a sentence expressing that there is an element containing exactly the atoms at an even position (in the ordering induced on the atoms) and containing the last atom. \square

Exercise 3.5.3 In the finite, $L_{\omega_1\omega}$ has the Beth property, the interpolation property, and is closed under order-invariant sentences. Hint: Use the fact that every class of finite structures is axiomatizable in $L_{\omega_1\omega}$ (cf. 3.2.1(b)). E.g.,

for the interpolation property argue as follows: If $\varphi \models_{\text{fin}} \psi$ for the $L_{\omega_1\omega}[\sigma]$ -sentence φ and the $L_{\omega_1\omega}[\tau]$ -sentence ψ then every $L_{\omega_1\omega}[\sigma \cap \tau]$ -sentence χ axiomatizing the class

$$\{\mathcal{A} \mid (\sigma \cap \tau) \mid \mathcal{A} \text{ a finite } \sigma\text{-structure with } \mathcal{A} \models \varphi\}$$

is an interpolant. □

Exercise 3.5.4 (a) For $\tau = \{<\}$ let φ be a $\{<\}$ -sentence expressing that $<$ is an ordering and, for $n \geq 1$, let χ_n be the sentence of FO^2 introduced in part (a) of 3.3.1, which says that an ordering has exactly n elements. Then (for both \models and \models_{fin}),

$$\models_{(\text{fin})} (\varphi(<) \wedge \bigvee_{n \geq 1} \chi_{2 \cdot n}(<)) \rightarrow (\varphi(<') \rightarrow \bigvee_{n \geq 1} \chi_{2 \cdot n}(<')).$$

Show that there is no interpolant in $L_{\infty\omega}^\omega$. Hence, $L_{\infty\omega}^\omega$ does not have the interpolation property (in the finite).

(b) For $s \geq 2$ and unary P , let $\varphi(P)$ and $\psi(P)$ be the FO^s sentences

$$\begin{aligned} \varphi(P) &:= \text{“there is exactly one element satisfying } P\text{”} \\ \psi(P) &:= \text{“there are at least } s \text{ elements not satisfying } P\text{”}. \end{aligned}$$

Then for unary Q ,

$$\models_{(\text{fin})} (\varphi(P) \wedge \psi(P)) \rightarrow (\varphi(Q) \rightarrow \psi(Q)),$$

but there is no interpolant in FO^s . Hence, FO^s does not have the interpolation property (in the finite).

(c) Show that FO^1 and $L_{\infty\omega}^1$ have the interpolation property (in the finite) in case τ is relational. □

Preservation Theorems

In the model theory for arbitrary structures certain closure properties of the class of models of a sentence φ are reflected by syntactic properties of φ . Most of these so-called preservation theorems fail when literally translated to the finite. We give some examples.

Call a first-order formula *universal* (*existential*) if it is built up from atomic and negated atomic formulas using only the connectives \wedge, \vee and the universal (existential) quantifier. If φ is a universal sentence, a simple inductive proof shows

$$B \subseteq A \text{ and } A \models \varphi \text{ imply } B \models \varphi,$$

that is, φ is *preserved under substructures*. Similarly, if φ is existential then

$$B \subseteq A \text{ and } B \models \varphi \text{ imply } A \models \varphi,$$

that is, φ is *preserved under extensions*. Moreover, in classical model theory one proves that every FO-sentence preserved under substructures is logically equivalent to a universal FO-sentence. We give an example of an FO-sentence that, in the finite, is preserved under substructures but is not equivalent to a universal first-order sentence.

Let the universal sentence φ_0 be the conjunction of the ordering axioms in $\{<, \min, \max\}$ and of the following sentence expressing that R is a “partial successor relation”,

$$\forall x \forall y (Rxy \rightarrow x < y) \wedge \forall x \forall y \forall z ((Rxy \wedge x < z) \rightarrow (y = z \vee y < z)).$$

Let φ_1 be the sentence $\forall x (\neg x = \max \rightarrow \exists y Rxy)$ expressing that R is the “total” successor relation. For finite structures \mathcal{A} and \mathcal{B} ,

$$(*) \quad \mathcal{A} \models \varphi_0, \mathcal{B} \models (\varphi_0 \wedge \varphi_1), \text{ and } \mathcal{B} \subseteq \mathcal{A} \text{ imply } \mathcal{A} = \mathcal{B}.$$

Using a new unary relation symbol Q , we set

$$\varphi := \varphi_0 \wedge (\varphi_1 \rightarrow \exists y Qy).$$

In finite models, φ is preserved under substructures: Suppose $(\mathcal{A}, Q^{\mathcal{A}}) \models \varphi$ and $(\mathcal{B}, Q^{\mathcal{B}}) \subseteq (\mathcal{A}, Q^{\mathcal{A}})$. Since φ_0 is universal, $\mathcal{B} \models \varphi_0$. If $\mathcal{B} \not\models \varphi_1$ then $(\mathcal{B}, Q^{\mathcal{B}}) \not\models \varphi$. If $\mathcal{B} \models \varphi_1$ then by (*), $\mathcal{B} = \mathcal{A}$, therefore $(\mathcal{B}, Q^{\mathcal{B}}) = (\mathcal{A}, Q^{\mathcal{A}})$ and hence, $(\mathcal{B}, Q^{\mathcal{B}}) \models \varphi$.

Assume by contradiction, that ψ is a universal first-order sentence with $\models_{\text{fin}} \varphi \leftrightarrow \psi$, say, $\psi = \forall x_1 \dots \forall x_n \chi$ with quantifier-free χ (every universal FO-formula is logically equivalent to one of this form!). Look at a $\{<, \min, \max, R\}$ -structure \mathcal{A} with $n + 3$ elements, where $(A, <^A, \min^A, \max^A)$ is an ordering and where R^A is the successor relation. Set $Q^{\mathcal{A}} = \emptyset$. Then $(\mathcal{A}, Q^{\mathcal{A}}) \not\models \varphi$, and hence, $(\mathcal{A}, Q^{\mathcal{A}}) \models \exists x_1 \dots \exists x_n \neg \chi$, say,

$$(\mathcal{A}, Q^{\mathcal{A}}) \models \neg \chi[a_1, \dots, a_n].$$

Choose $a \in A \setminus \{a_1, \dots, a_n, \min^A, \max^A\}$ and set $Q' = \{a\}$. Since χ is quantifier-free, $(\mathcal{A}, Q') \models \neg \chi[a_1, \dots, a_n]$ and therefore, $(\mathcal{A}, Q') \not\models \forall x_1 \dots \forall x_n \chi$. On the other hand, $(\mathcal{A}, Q') \models \varphi$. Hence φ and ψ are not equivalent in the finite.

This example gives part (a) of

Proposition 3.5.5 (a) *There is a first-order sentence which, in the finite, is preserved under substructures but not equivalent to a universal first-order sentence.*

(b) *There is a first-order sentence which, in the finite, is preserved under extensions but not equivalent to an existential first-order sentence.*

Proof. (b) Let φ be according to (a). Then $\neg \varphi$ is preserved under extensions and not equivalent to an existential sentence. \square

We close with some remarks concerning “monotone” formulas. Fix a relation symbol R of τ of arity r . A sentence φ is *monotone in R (in the finite)* if

$$(\mathcal{A}, R_1) \models \varphi, (\mathcal{A} \text{ finite}) \text{ and } R_1 \subseteq R_2 \subseteq A^r \text{ imply } (\mathcal{A}, R_2) \models \varphi.$$

A first-order formula φ is *positive* in R if φ is built up from atomic formulas using $\neg, \wedge, \vee, \forall, \exists$ and any occurrence of the relation symbol R in φ is within the scope of an even number of negation symbols. An easy inductive argument shows that a sentence positive in R is monotone. While any first-order sentence monotone in R is logically equivalent to a formula positive in R , this is no longer true in the finite. We state this result; for a proof we refer to the example in Exercise 3.5.8.

Proposition 3.5.6 *There is a first-order sentence which, in the finite, is monotone in R , but not equivalent to a first-order sentence positive in R .*

As we have seen, preservation theorems of first-order logic fail when literally translated to the case of finite structures. However, it is still open whether – say, in the case of sentences preserved under substructures – there is a syntactically defined or, at least, recursive set of sentences preserved under substructures in the finite such that, in the finite, every first-order sentence preserved under substructures is equivalent to a sentence in this set.

Exercise 3.5.7 Define the notions of a universal $L_{\omega_1\omega}$ -sentence and of an $L_{\omega_1\omega}$ -sentence positive in R and show:

- (a) In the finite, every $L_{\omega_1\omega}$ -sentence preserved under substructures is equivalent to a universal $L_{\omega_1\omega}$ -sentence.
- (b) In the finite, every $L_{\omega_1\omega}$ -sentence monotone in R is equivalent to an $L_{\omega_1\omega}$ -sentence positive in R .

Hint for (a): Consider a variant of the Ehrenfeucht game by prescribing that in each move the spoiler has to choose an element of the second structure and argue with the corresponding isomorphism types defined by

$$\chi_{\bar{a}}^0 := \varphi_{\bar{a}}^0 \quad \text{and} \quad \chi_{\bar{a}}^{j+1} := \forall v \bigvee_{a \in A} \chi_{\bar{a}a}^j.$$

For (b) omit subformulas $\neg R t_1 \dots t_m$ in $\varphi_{\bar{a}}^0$, defining $\varphi_{\bar{a}}^m$ for $m > 0$ as usual. □

Exercise 3.5.8 (a) Assume that $\tau = \sigma \cup \{R\}$, where $R \notin \sigma$ is an r -ary relation symbol, and let \mathcal{A} and \mathcal{B} be τ -structures. Let $\text{Part}^+(\mathcal{A}, \mathcal{B})$ be the set

$$\{p \in \text{Part}(\mathcal{A}|\sigma, \mathcal{B}|\sigma) \mid \text{for all } \bar{a} \in \text{do}(p), \text{ if } R^{\mathcal{A}}\bar{a} \text{ then } R^{\mathcal{B}}p(\bar{a})\}.$$

Write

$$\mathcal{A} \equiv_m^+ \mathcal{B} \quad \text{iff} \quad \text{for all FO}[\tau]\text{-sentences } \varphi \text{ with } \text{qr}(\varphi) \leq m \text{ that are positive in } R: \text{ if } \mathcal{A} \models \varphi \text{ then } \mathcal{B} \models \varphi,$$

and

$\mathcal{A} \cong_m^+ \mathcal{B}$ iff there is a sequence $(I_j)_{j \leq m}$ of nonempty subsets of $\text{Part}^+(\mathcal{A}, \mathcal{B})$ with the back and forth properties.

Show that $\mathcal{A} \equiv_m^+ \mathcal{B}$ iff $\mathcal{A} \cong_m^+ \mathcal{B}$.

(b) Let $\sigma = \{S\}$ and $\tau = \sigma \cup \{\prec\}$ with binary S, \prec . For $n, k \in \mathbb{N}$ let $\mathcal{G}_{n,k}$ be the structure with

$$\begin{aligned} G_{n,k} &:= \{0, \dots, n\} \times \{0, \dots, k\}, \\ S^{G_{n,k}} &:= \{((i, j), (i, j+1)) \mid i \leq n, j < k\} \\ &\quad \cup \{((i, j), (i+1, j+1)) \mid i < n, j < k\} \\ &\quad \cup \{((n, j), (0, j+1)) \mid j < k\}, \\ \prec^{G_{n,k}} &:= \text{transitive closure of } S^{G_{n,k}}. \end{aligned}$$

$\mathcal{G}_{n,k}$ can be pictured as a circular net of circumference $(n+1)$ and height k , where S -arrows are going from a point (that is not in the upper edge) to its upper and to its upper right neighbour. We say that a point (i, j) has (geographical) length i and (geographical) latitude j .

If $a \prec^{G_{n,k}} b$, let $\mathcal{G}_{n,k}^{(a,b)}$ be defined as $\mathcal{G}_{n,k}$, setting $\prec^{G_{n,k}^{(a,b)}} := \prec^{G_{n,k}} \setminus \{(a, b)\}$.

Show for $n = 4 \cdot (2^m + 1)$: If $a \prec^{G_{n,n}} b$, a has latitude $2^m + 1$, b has latitude $3 \cdot (2^m + 1)$, and if a and b have length difference $k := 2 \cdot (2^m + 1)$, then

$$\mathcal{G}_{n,k} \cong_m^+ \mathcal{G}_{n,n}^{(a,b)}.$$

Hint: Let c and d in $G_{n,k}$ have latitude $2^m + 1$ and the lengths of a and b , respectively, and show that $(\mathcal{G}_{n,k}, c, d) \cong_m^+ (\mathcal{G}_{n,n}^{(a,b)}, a, b)$. For $l \leq m$, take as I_l the set of those $p \in \text{Part}^+(\mathcal{G}_{n,k}, c, d), (\mathcal{G}_{n,n}^{(a,b)}, a, b)$ that

- preserve lengths;
- preserve the upper and the lower edge;
- preserve differences of latitudes up to 2^l for points whose lengths differ by at most 2^l and for points on the edges;
- do not decrease differences of latitudes.

(c) Set $\varphi = \forall x \forall y (Sxy \rightarrow (\neg x = y \wedge x \prec y)) \wedge (\varphi_1 \vee \varphi_2 \vee \varphi_3)$ where

$$\begin{aligned} \varphi_1 &:= \text{“} \prec \text{ is transitive”} \\ \varphi_2 &:= \exists x \exists y (x \prec y \wedge y \prec x) \\ \varphi_3 &:= \exists x \exists y (x \prec y \wedge \forall z (Szy \rightarrow (\neg z \preceq z \wedge \kappa(z)))) \end{aligned}$$

with

$$\begin{aligned} \kappa(z) &:= \forall w \forall y \forall u \forall v ((y \preceq w \wedge Swz \wedge v \preceq u \wedge Suy) \rightarrow \\ &\quad (v \prec y \wedge v \prec w \wedge v \prec z \wedge y \prec z)). \end{aligned}$$

$\kappa(z)$ guarantees that \prec is transitive on the set consisting of z and its iterated predecessors; more precisely, if $\mathcal{A} \models \kappa[a_0]$ and $S^A a_{i+1} a_i$ for $i \geq 0$ then $a_i \prec a_j$ for $0 \leq j < i$.

Show that $\mathcal{G}_{n,k} \models \varphi$ and $\mathcal{G}_{n,n}^{(a,b)} \models \neg\varphi$, where n, k, a, b are as in (b).

(d) Show that, in the finite, φ is monotone in \prec but not equivalent to a sentence positive in \prec . Hint for the monotonicity: Let \mathcal{A}, \mathcal{B} be given such that $B = A$, $S^B = S^A$, and $\prec^B \supseteq \prec^A$ and such that $\mathcal{A} \models \varphi$. As φ_2 is positive in \prec , assume $\mathcal{A} \models \neg\varphi_2$. Consider the cases $\mathcal{A} \models \varphi_3$ and $\mathcal{A} \models \varphi_1 \wedge \neg\varphi_3$. In the first case, if (a, b) witnesses φ_3 in \mathcal{A} but not in \mathcal{B} , find a chain $(a, b) = (a_0, b_0), (a_1, b_1), \dots$ with $a_i \prec^B b_i, (b_{i+1}, b_i) \in \text{TC}(S^A)$, and $(a_i, b_i) \notin \text{TC}(S^A)$ that ends with a pair witnessing φ_3 in \mathcal{B} .

(e) Show that φ is not monotone in \prec for arbitrary structures. Hint: Start with a model of φ that consists of the negative integers with the natural successor relation and ordering together with an additional point. \square

Notes 3.5.9 A reference for second-order logic is [133]; for the model theory of $L_{\infty\omega}$ and $L_{\omega_1\omega}$ we refer the reader to [8, 100]. The first study of fragments with finitely many variables is due to Henkin [84]. The languages $L_{\infty\omega}^s$ and the corresponding pebble games were first introduced by Barwise [10] and then reinvented and successfully used by Immerman [90] and Poizat [130]. Further references are [96] and [40]. Compare [30] for most results in subsection 3.3.3. An exposition of the model theory of $L_{\infty\omega}^\omega$ is contained in [86]. Surveys on recent results for finite variable logics are [58, 66]. (Infinite) counting quantifiers and the corresponding Ehrenfeucht-Fraïssé method have been considered in various contexts, e.g. in [145]. A thorough treatment in the context of finite model theory is contained in [125]. The results on graphs presented in section 3.4 are from [97]. The failure of many classical interpolation and preservation theorems was first presented explicitly in [70]. The example given in 3.5.8 is from [135], the first such example was presented in [6].

4. 0–1 Laws

So far we mainly looked at techniques and results which were developed for *arbitrary* structures, and we analyzed to what extent they remain valid if we restrict ourselves to *finite* ones. In the present chapter we study a concept that is genuine to the finite (even though there are extensions to arbitrary structures): the probability for a finite structure to be a model of a given sentence. A starting point was the observation that for a relational vocabulary τ and a first-order sentence φ , either nearly all finite structures are models of φ or nearly all finite structures are models of $\neg\varphi$ (for example, nearly all $\{E\}$ -structures are not digraphs, that is, nearly all $\{E\}$ -structures are models of $\neg\forall x\neg Exx$). One describes this property by saying that first-order logic satisfies the 0–1 law.

There are different ways of how to count structures and hence, there are different versions of 0–1 laws which we study for FO and for $L_{\infty\omega}^\omega$ in the first and the third section. The second section treats relativized versions, for example, 0–1 laws inside a class of structures, say, inside the class of graphs. After some applications in section 4.4, we study the case of monadic second-order logic and disprove the 0–1 law for it in section 4.5.

In this chapter n is always a natural number greater than 0.

4.1 0–1 Laws for FO and $L_{\infty\omega}^\omega$

As always we refer to a fixed vocabulary τ . For a class K of structures let $L_n(K)$ be the number of structures in K with universe $\{1, 2, \dots, n\}$,

$$L_n(K) := |\{A \in K \mid A = \{1, \dots, n\}\}|.$$

Sometimes, structures A with $A = \{1, \dots, n\}$ are called *labeled* structures, since every element in such a structure is labeled with a natural number. Hence, $L_n(K)$ is the number of labeled structures in K of cardinality n . If K is the class of models of a sentence φ or the class of all τ -structures, we denote $L_n(K)$ also by $L_n(\varphi)$ and $L_n(\tau)$, respectively. Let $l_n(K)$ be the fraction of structures with universe $\{1, \dots, n\}$ which are in K ,

$$l_n(K) := \frac{L_n(K)}{L_n(\tau)}.$$

In case it exists, $l(K) := \lim_{n \rightarrow \infty} l_n(K)$ is called the *labeled asymptotic probability* of K . Similarly as above, $l_n(\varphi)$ stands for $l_n(\text{Mod}(\varphi))$ and $l(\varphi)$ for $l(\text{Mod}(\varphi))$. If $l(\varphi) = 1$ we say that φ holds in *almost all* finite structures or that φ *almost surely* holds.

A class Ψ of sentences of a logic is said to satisfy the *labeled 0–1 law* if

$$l(\varphi) = 1 \quad \text{or} \quad l(\varphi) = 0$$

holds for every $\varphi \in \Psi$ (or, equivalently, for $\varphi \in \Psi$ either φ or $\neg\varphi$ holds in almost all finite structures).

Examples 4.1.1 (a) Suppose $\tau = \{P, c\}$, where P is unary. Since for any τ -structure (A, P^A, c^A) ,

$$(A, P^A, c^A) \models Pc \quad \text{iff} \quad (A, A \setminus P^A, c^A) \not\models Pc,$$

we see that $l_n(Pc) = \frac{1}{2}$ and thus, $l(Pc) = \frac{1}{2}$.

(b) For the “vocabulary” $\tau := \{f\}$ with a unary function symbol f consider the first-order sentence $\forall x f(x) \neq x$ expressing that f has no fixed-point. Then

$$l_n(\forall x f(x) \neq x) = \left(\frac{n-1}{n}\right)^n = \left(1 - \frac{1}{n}\right)^n$$

(since on the universe $\{1, \dots, n\}$ one can fix the values of f independently and since for each argument i the $n-1$ possible values $\neq i$ do not lead to a fixed-point). Hence, $l(\forall x f(x) \neq x) = e^{-1}$.

(c) If K is the class $\text{EVEN}[\tau]$ of all τ -structures of even finite cardinality, then

$$l_n(K) = \begin{cases} 1 & \text{if } n \text{ is even} \\ 0 & \text{if } n \text{ is odd} \end{cases}$$

and hence, $l(K)$ does not exist. Therefore, $l(\varphi)$ does not exist for the second-order sentence φ expressing “there is a binary relation which is an equivalence relation having only equivalence classes with exactly two elements” or for the $L_{\omega_1\omega}$ -sentence $\bigvee_{k \geq 1} \varphi_{=2k}$.

(d) Let K be a class of structures for a relational τ . Construct a “random” structure of vocabulary τ on $\{1, \dots, n\}$ by the following experiment: For every m -ary relation symbol R in τ and for every $i_1, \dots, i_m \in \{1, \dots, n\}$, toss a fair coin to decide whether $Ri_1 \dots i_m$ is true. Then $l_n(K)$ is the probability for the outcome \mathcal{A} of the experiment to belong to K . \square

Examples (a) and (b) show that we can expect a labeled 0–1 law for first-order logic only for relational vocabularies. In fact, it then holds. The central point of the proof we give is the fact that the extension axioms hold almost surely and that for any φ , either φ or $\neg\varphi$ is a consequence of the extension axioms.

In the following suppose that τ is relational. Recall (cf. 3.2.11) that an $(r+1)$ -extension axiom is a sentence $\chi_{\Phi} =$

$$\forall v_1 \dots \forall v_r \left(\bigwedge_{1 \leq i < j \leq r} v_i \neq v_j \rightarrow \exists v_{r+1} \left(\bigwedge_{1 \leq i \leq r} v_i \neq v_{r+1} \wedge \bigwedge_{\varphi \in \Phi} \varphi \wedge \bigwedge_{\varphi \in \Phi^c} \neg \varphi \right) \right)$$

where Φ is a subset of

$$\Delta_{r+1} = \{ \varphi(v_1, \dots, v_r, v_{r+1}) \mid \varphi \text{ has the form } R\bar{x}, \text{ where } R \in \tau \text{ and where } v_{r+1} \text{ occurs in } \bar{x} \}.$$

Lemma 4.1.2 *Any extension axiom holds in almost all finite structures.*

Proof. Given Φ , we have to show that the asymptotic labeled probability $l(\chi_\Phi)$ equals 1. For any tuple a_1, \dots, a_r of distinct elements in a structure \mathcal{A} and any further object a let δ be the probability that a_1, \dots, a_r, a satisfies $\Phi \cup \{ \neg \varphi \mid \varphi \in \Phi^c \}$, when adding a to \mathcal{A} as a new element and randomly fixing the truth values of $R\bar{b}$ for any R in τ and any sequence \bar{b} in $A \cup \{a\}$ containing a . Clearly, if c is the number of subsets of Δ_{r+1} , then $\delta = \frac{1}{c}$; in particular, $\delta > 0$. Thus

$$\begin{aligned} l_n(\neg \chi_\Phi) &= l_n(\exists v_1 \dots \exists v_r (\bigwedge_{1 \leq i < j \leq r} v_i \neq v_j \wedge \\ &\quad \forall v_{r+1} (\bigvee_{1 \leq i \leq r} v_i = v_{r+1} \vee \bigvee_{\varphi \in \Phi} \neg \varphi \vee \bigvee_{\varphi \in \Phi^c} \varphi))) \\ &\leq n^r \left(\frac{c-1}{c} \right)^{n-r} = n^r (1 - \delta)^{n-r}. \end{aligned}$$

Therefore, $l(\neg \chi_\Phi) = \lim_{n \rightarrow \infty} l_n(\neg \chi_\Phi) = 0$. □

Recall that by T_{rand} ($= T_{\text{rand}}(\tau)$) we denote the set of extension axioms.

Corollary 4.1.3 *Let φ be a first-order sentence.*

(a) *If $T_{\text{rand}} \models \varphi$ then $l(\varphi) = 1$.* (b) *If $T_{\text{rand}} \models \neg \varphi$ then $l(\varphi) = 0$.*

Proof. (a) If $T_{\text{rand}} \models \varphi$ then by compactness, $T_0 \models \varphi$ for some finite subset T_0 of T_{rand} . Since T_0 is a set of extension axioms, $l(\bigwedge T_0) = 1$ by the preceding lemma. Hence, $l(\varphi) = 1$.

(b) If $T_{\text{rand}} \models \neg \varphi$ then by (a), $l(\neg \varphi) = 1$ and therefore, $l(\varphi) = 0$. □

For $s \geq 1$, let ϵ_s be the conjunction of the finitely many r -extension axioms with $r \leq s$. Since $l(\epsilon_s) = 1$ we obtain

Corollary 4.1.4 *Let φ be an $L_{\infty\omega}^\omega$ -sentence.*

(a) *If $\epsilon_s \models \varphi$ then $l(\varphi) = 1$.* (b) *If $\epsilon_s \models \neg \varphi$ then $l(\varphi) = 0$.* □

Theorem 4.1.5 *Let τ be relational. Then both $\text{FO}[\tau]$ and $L_{\infty\omega}^\omega[\tau]$ satisfy the labeled 0–1 law.*

Proof. The assertions follow from the above corollaries, since we know that for φ in $\text{FO}[\tau]$ (cf. 3.2.11)

$$T_{\text{rand}} \models \varphi \quad \text{or} \quad T_{\text{rand}} \models \neg \varphi$$

and that for $\varphi \in L_{\infty\omega}^s[\tau]$ (cf. 3.3.11),

$$\epsilon_s \models \varphi \quad \text{or} \quad \epsilon_s \models \neg \varphi. \quad \square$$

Exercise 4.1.6 Let τ be relational and $m \in \mathbb{N}$. For $n, k \geq 1$ let $f(n, k)$ be the probability that $\mathcal{A} \cong_m \mathcal{B}$ holds for random structures \mathcal{A} and \mathcal{B} with universe $\{1, \dots, n\}$ and $\{1, \dots, k\}$, respectively. Show that $\lim_{n, k \rightarrow \infty} f(n, k) = 1$. \square

Exercise 4.1.7 Let $0 < p < 1$. Alter the construction of 4.1.1(d) of a random structure with universe $\{1, \dots, n\}$ using a biased coin, so that $Ri_1 \dots i_m$ now holds with probability p . For any extension axiom φ show that with probability one this random structure is a model of φ . The situation changes drastically if p is allowed to depend on the number n of elements (see [132]). \square

Exercise 4.1.8 Let τ be relational. Extending the construction of 4.1.1(d), define an infinite random structure over $\{1, 2, \dots\}$ in the following way: For every P in τ , say of arity m , and every m -tuple i_1, \dots, i_m of positive integers toss a coin to decide whether $Pi_1 \dots i_m$ is true. Show that with probability 1 (with respect to the canonical product measure) this infinite random structure is a model of T_{rand} and hence, it is the (up to isomorphism) uniquely determined countable model \mathcal{R} of T_{rand} (cf. 3.2.11). Therefore, \mathcal{R} is called the *infinite random structure*. \square

Exercise 4.1.9 Show that for the infinite “vocabulary” $\tau = \{P_1, P_2, \dots\}$ with unary P_i there is an $L_{\omega_1\omega}^\omega[\tau]$ -sentence without labeled asymptotic probability. \square

4.2 Parametric Classes

Does every first-order sentence or its negation hold in almost all graphs? To treat such questions we introduce conditional probabilities. Slight modifications of the arguments given in the preceding section show that the 0–1 law holds for the conditional probabilities with respect to classes of structures which are axiomatizable by means of “parametric” axioms. The classes of graphs and digraphs are prominent examples.

We start with the definition of conditional probabilities: Suppose that K and H are classes of τ -structures. Define the *labeled probability* $l_n(K|H)$ by

$$l_n(K|H) := \frac{L_n(K \cap H)}{L_n(H)}.^1$$

If it exists, $l(K|H) := \lim_{n \rightarrow \infty} l_n(K|H)$ is called the *labeled asymptotic probability* of K with respect to H . Notations such as $l_n(\varphi|H)$ or $l_n(K|\tau)$ should be self-explaining.

Obviously, $l_n(K|\tau) = l_n(K)$, and $l_n(\text{CONN}|\text{GRAPH})$ is the number of connected graphs on $\{1, \dots, n\}$ divided by the total number of graphs on $\{1, \dots, n\}$.

¹ This fraction is defined only in case H contains a structure of cardinality n .

Let \forall distinct $x_1 \dots x_s \psi$ abbreviate $\forall x_1 \dots \forall x_s ((\neg x_1 = x_2 \wedge \neg x_1 = x_3 \wedge \dots \wedge \neg x_{s-1} = x_s) \rightarrow \psi)$ (which is $\forall x_1 \psi$ in case $s = 1$).

Definition 4.2.1 Let τ be relational. A first-order sentence φ is called *parametric*, if it is a conjunction of sentences \forall distinct $x_1 \dots x_s \psi$, where $s \geq 1$ and ψ is a boolean combination of formulas of the form $Ry_1 \dots y_n$ with $R \in \tau$ and $\{y_1, \dots, y_n\} = \{x_1, \dots, x_s\}$. (Note that s cannot exceed the maximum of the arities of relation symbols in τ .) A class K of structures is said to be *parametric*, if $K = \text{Mod}(\varphi)$ where φ is a parametric sentence. \square

Examples 4.2.2 (a) $\forall x \neg Rxx \wedge \forall$ distinct $xy (Rxy \rightarrow Ryx)$ and $\forall x \neg Rxx$ are parametric sentences axiomatizing the classes of graphs and digraphs, respectively.

(b) The class of *tournaments* is axiomatized by the parametric sentence

$$\forall x \neg Rxx \wedge \forall \text{ distinct } xy (Rxy \leftrightarrow \neg Ryx).$$

(c) For relational τ the class K of all τ -structures is parametric, since K is the class of models of the empty conjunction.

(d) Note that

$$\forall \text{ distinct } xyz ((Rxy \wedge Ryz) \rightarrow Rxz)$$

is not a parametric sentence (since, e.g., $\{x, y\} \neq \{x, y, z\}$). In fact, from our analysis of parametric classes it will become clear that, for example, the classes of transitive relations, equivalence relations, partial orderings, and orderings are not parametric, transitivity being the only obstacle.

(e) If R is k -ary, \forall distinct $x_1 \dots x_k (Rx_1 \dots x_k \wedge \neg Rx_1 \dots x_k)$ is a parametric sentence that is true in all structures of cardinality $< k$, but has no model of cardinality $\geq k$. \square

The following considerations will demonstrate that the sentence in (e) is a “sharpest” example of a parametric sentence having only models of bounded cardinality.

Let k be the maximum of the arities of the relation symbols in τ . A parametric sentence (and its model class) is called *nontrivial*, if it has a model of cardinality $\geq k$. We show that nontrivial parametric sentences have arbitrarily large models: Suppose that φ_0 is a nontrivial parametric sentence. If B is any nonempty set, then the following procedure, where we stepwise fix the relations on B , leads to a model of φ_0 with universe B : For any $s \leq k$ and any distinct $b_1, \dots, b_s \in B$, we choose an arbitrary model \mathcal{A} of φ_0 of cardinality $\geq s$ and distinct $a_1, \dots, a_s \in A$ and define, for all R in τ , the “ $\{b_1, \dots, b_s\}$ -part of R^B ” as a copy of the “ $\{a_1, \dots, a_s\}$ -part of R^A ”; more precisely, we define this part of R^B such that for $\varphi(v_1, \dots, v_s) = Ry_1 \dots y_n$ with $\{y_1, \dots, y_n\} = \{v_1, \dots, v_s\}$

$$\mathcal{B} \models \varphi[b_1, \dots, b_s] \quad \text{iff} \quad \mathcal{A} \models \varphi[a_1, \dots, a_s].$$

Since in \mathcal{B} every s -tuple of distinct elements behaves as some s -tuple in some model of φ_0 and as φ_0 is parametric, \mathcal{B} itself is a model of φ_0 . Thus we have proved:

(1) φ_0 has models in all cardinalities.

Call an $(r+1)$ -extension axiom

$$\forall \text{ distinct } v_1 \dots v_r \exists v_{r+1} \left(\bigwedge_{1 \leq i \leq r} \neg v_i = v_{r+1} \wedge \bigwedge_{\varphi \in \Phi} \varphi \wedge \bigwedge_{\varphi \in \Phi^c} \neg \varphi \right)$$

compatible with φ_0 if

$$\{\varphi_0\} \cup \{ \exists v_1 \dots \exists v_r \exists v_{r+1} \left(\bigwedge_{1 \leq i < j \leq r+1} \neg v_i = v_j \wedge \bigwedge_{\varphi \in \Phi} \varphi \wedge \bigwedge_{\varphi \in \Phi^c} \neg \varphi \right) \}$$

is satisfiable. Let $T_{\text{rand}}(\varphi_0)$ be the set of sentences consisting of φ_0 and of all extension axioms compatible with φ_0 . The proof in Example 3.2.11 showing that any two models of T_{rand} are partially isomorphic, also works for $T_{\text{rand}}(\varphi_0)$ and yields:

(2) Any two models of $T_{\text{rand}}(\varphi_0)$ are partially isomorphic and hence, $L_{\infty\omega}$ -equivalent. Therefore,

$$T_{\text{rand}}(\varphi_0) \models \psi \quad \text{or} \quad T_{\text{rand}}(\varphi_0) \models \neg\psi$$

holds for any $L_{\infty\omega}$ -sentence ψ .

Also the proof of 3.2.11 leading to a countable model of T_{rand} can be transferred to $T_{\text{rand}}(\varphi_0)$, since by the construction process described above the corresponding \mathcal{A}_i 's can be chosen as models of φ_0 . Hence,

(3) $T_{\text{rand}}(\varphi_0)$ has an (up to isomorphism) unique countable model $\mathcal{R}(\varphi_0)$.

For $s \geq 1$, we denote by φ_0^s the conjunction of φ_0 with the finitely many r -extension axioms with $r \leq s$ that are compatible with φ_0 . By a similar argumentation, one obtains:

(4) Any two models of φ_0^s are s -partially isomorphic and hence, $L_{\infty\omega}^s$ -equivalent. Therefore, for any $L_{\infty\omega}^s$ -sentence ψ ,

$$\varphi_0^s \models \psi \quad \text{or} \quad \varphi_0^s \models \neg\psi.$$

Finally, we have:

(5) If ψ is an extension axiom compatible with φ_0 then $l(\psi|\varphi_0) = 1$.

To give a proof, one can argue as for 4.1.2, but now restricting oneself to models of φ_0 and taking as c the number of subsets Φ of Δ_{r+1} that correspond to $(r+1)$ -extension axioms compatible with φ_0 . (Given Φ , distinct a_1, \dots, a_r ,

and a new a , we can satisfy Φ by $a_1 \dots a_r a$, applying the construction procedure described above to any $a_{i_1} \dots a_{i_m} a$ with $1 \leq i_1 < i_2 < \dots < i_m \leq r$.)

Let H be a class of structures and Ψ a class of sentences. We say that H satisfies the labeled 0–1 law for Ψ if

$$l(\psi|H) = 1 \quad \text{or} \quad l(\psi|H) = 0$$

holds for any $\psi \in \Psi$. From (2)–(5) we obtain:

Theorem 4.2.3 *Let H be a nontrivial parametric class. Then H satisfies the labeled 0–1 law for $L_{\infty\omega}^\omega$ and hence, for FO.* □

4.3 Unlabeled 0–1 Laws

We saw in 4.1.1(d) that from a probabilistic point of view the definition of $l_n(K)$ is quite natural. But note that for $\tau := \{P\}$ with unary P and $i = 1, \dots, n$ the structures $\mathcal{A}_i := (\{1, \dots, n\}, P^{A_i})$ with $P^{A_i} := \{i\}$ are counted as n different structures in the definition of $l_n(K)$, even though they are isomorphic. In this section we study the so-called *unlabeled probability* $u_n(K)$, which is the proportion of *isomorphism types* of structures of cardinality n in K . Similarly, we define unlabeled conditional probabilities. It turns out that the labeled and unlabeled asymptotic conditional probabilities coincide in case almost all structures in the underlying class are rigid.

Fix a vocabulary τ . For a class K let $U_n(K)$ be the number of isomorphism types of structures of cardinality n in K or, equivalently,

$$U_n(K) \quad := \quad \text{number of isomorphism types of} \\ \text{structures in } K \text{ with universe } \{1, \dots, n\}.$$

$U_n(\tau)$ and $U_n(\varphi)$ stand for $U_n(K)$, where K is the class of all τ -structures and the class of models of φ , respectively.

For arbitrary K we set

$$u_n(K) \quad := \quad \frac{U_n(K)}{U_n(\tau)}$$

and denote (in case it exists) by $u(K) := \lim_{n \rightarrow \infty} u_n(K)$ the *unlabeled asymptotic probability*.

If K and H are classes of structures we call

$$u_n(K|H) \quad := \quad \frac{U_n(K \cap H)}{U_n(H)}$$

the *unlabeled probability* and

$$u(K|H) := \lim_{n \rightarrow \infty} u_n(K|H)$$

the *unlabeled asymptotic probability* of K with respect to H .

Thus, $u_n(\text{CONN}|\text{GRAPH})$ is the number of isomorphism types of connected n vertex graphs divided by the total number of isomorphism types of n vertex graphs.

When studying the relationship between labeled and unlabeled probabilities, a prominent role is played by the class RIG ($= \text{RIG}[\tau]$) of rigid structures. A structure \mathcal{A} is called *rigid* if the identity on \mathcal{A} is the only automorphism of \mathcal{A} .

- Lemma 4.3.1** (a) $L_n(K) \leq U_n(K) \cdot n!$
 (b) If $K \subseteq \text{RIG}$ then $L_n(K) = U_n(K) \cdot n!$
 (c) If $K \subseteq \text{RIG}^c$ ² then $L_n(K) \leq U_n(K) \cdot \frac{n!}{2}$.

Proof. Since $L_n(K) = L_n(K \cap \text{RIG}) + L_n(K \cap \text{RIG}^c)$ and similarly for U_n , it suffices to show (b) and (c).

Let \mathcal{A} be a structure with universe $\{1, \dots, n\}$. There are $n!$ permutations of $\{1, \dots, n\}$. Every permutation π gives a structure \mathcal{A}_π on $\{1, \dots, n\}$ such that $\pi : \mathcal{A} \cong \mathcal{A}_\pi$. Clearly, for permutations π and ρ of $\{1, \dots, n\}$ we have

$$(*) \quad \mathcal{A}_\pi = \mathcal{A}_\rho \quad \text{iff} \quad \pi^{-1} \circ \rho : \mathcal{A} \cong \mathcal{A}.$$

Hence, if \mathcal{A} is rigid, we have

$$\mathcal{A}_\pi = \mathcal{A}_\rho \quad \text{iff} \quad \pi = \rho.$$

Thus each rigid structure leads to $n!$ distinct structures on $\{1, \dots, n\}$. This shows (b).

If \mathcal{A} is not rigid and ρ is a nontrivial automorphism of \mathcal{A} , then by (*),

$$\mathcal{A}_\pi = \mathcal{A}_{\pi \circ \rho}$$

for any permutation π . Hence, any nonrigid structure leads to at most $\frac{n!}{2}$ distinct structures on $\{1, \dots, n\}$. This proves (c). □

Lemma 4.3.2 For any class H , $u_n(\text{RIG}|H) \leq l_n(\text{RIG}|H)$. In particular,

$$u(\text{RIG}|H) = 1 \quad \text{implies} \quad l(\text{RIG}|H) = 1.$$

Proof.

$$\begin{aligned} u_n(\text{RIG}|H) &= \frac{U_n(\text{RIG} \cap H) \cdot n!}{U_n(\text{RIG} \cap H) \cdot n! + U_n(\text{RIG}^c \cap H) \cdot n!} \\ &\leq \frac{L_n(\text{RIG} \cap H)}{L_n(\text{RIG} \cap H) + L_n(\text{RIG}^c \cap H)} \quad (\text{by 4.3.1}) \\ &= l_n(\text{RIG}|H). \end{aligned}$$

□

² Recall that RIG^c denotes the complement of RIG , the class of τ -structures not in RIG .

As suggested by (b) and (c) of Lemma 4.3.1, almost all structures in a class H are rigid iff $L_n(H) \approx U_n(H) \cdot n!$. In fact:

Proposition 4.3.3 *Let H be a class of structures. Then*

$$u(\text{RIG}|H) = 1 \quad \text{iff} \quad \lim_{n \rightarrow \infty} \frac{L_n(H)}{U_n(H) \cdot n!} = 1.$$

Proof. As

$$\begin{aligned} \frac{L_n(H)}{U_n(H) \cdot n!} &= \frac{L_n(\text{RIG} \cap H)}{U_n(H) \cdot n!} + \frac{L_n(\text{RIG}^c \cap H)}{U_n(H) \cdot n!} \\ &= u_n(\text{RIG}|H) + \frac{L_n(\text{RIG}^c \cap H)}{U_n(H) \cdot n!} \quad (\text{by 4.3.1(b)}), \end{aligned}$$

we get, using 4.3.1(c),

$$u_n(\text{RIG}|H) \leq \frac{L_n(H)}{U_n(H) \cdot n!} \leq u_n(\text{RIG}|H) + \frac{1}{2}u_n(\text{RIG}^c|H) = 1 - \frac{1}{2}u_n(\text{RIG}^c|H).$$

So, $u(\text{RIG}|H) = 1$ implies $\lim_{n \rightarrow \infty} \frac{L_n(H)}{U_n(H) \cdot n!} = 1$, and $\lim_{n \rightarrow \infty} \frac{L_n(H)}{U_n(H) \cdot n!} = 1$ implies $u(\text{RIG}^c|H) = 0$, that is, $u(\text{RIG}|H) = 1$. \square

The following theorem will be used to extend for parametric classes the 0-1 laws from the labeled to the unlabeled case.

Theorem 4.3.4 *Let H be a class of structures. If almost all structures in H are rigid, i.e., if $u(\text{RIG}|H) = 1$, then for any class K the labeled and the unlabeled asymptotic probabilities with respect to H coincide, that is,*

$$l(K|H) = u(K|H) \quad (= u(K|\text{RIG} \cap H)).^3$$

Proof. By assumption, $u(\text{RIG}|H) = 1$ and therefore, $l(\text{RIG}|H) = 1$ by 4.3.2. Immediately from the definitions we get

$$l(K|H) = l(K|\text{RIG} \cap H) \quad \text{and} \quad u(K|H) = u(K|\text{RIG} \cap H).$$

Therefore, the claim follows from

$$\begin{aligned} l_n(K|\text{RIG} \cap H) &= \frac{L_n(K \cap \text{RIG} \cap H)}{L_n(\text{RIG} \cap H)} = && (\text{by 4.3.1(b)}) \\ &= \frac{U_n(K \cap \text{RIG} \cap H) \cdot n!}{U_n(\text{RIG} \cap H) \cdot n!} = u_n(K|\text{RIG} \cap H). \end{aligned}$$

\square

³ Clearly, $\lim_{n \rightarrow \infty} l_n(K|H) = \lim_{n \rightarrow \infty} u_n(K|H)$ means that both sides converge to the same number or else that both of them diverge.

Let H_0 be the parametric class consisting of all τ -structures \mathcal{A} such that $R^{\mathcal{A}} = \emptyset$ for all $R \in \tau$. Then $L_n(H_0) = U_n(H_0) = 1$ and $u_n(\text{RIG}|H_0) = 0$ (for $n \geq 2$). On the other hand, almost all structures in a “free” parametric class are rigid. Intuitively speaking, a parametric class H is free if for some $r \geq 2$ there is a real choice when fixing the parts of the relations corresponding to r -tuples of distinct elements. More precisely: If $H = \text{Mod}(\varphi_0)$ with a parametric φ_0 , then H is *free*, if for some $m \geq 2$ there is a relation symbol R , say of arity r , and a surjection $i : \{1, \dots, r\} \rightarrow \{1, \dots, m\}$ such that

$$\varphi_0 \wedge \exists x_1 \dots \exists x_m (Rx_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} \neg x_k = x_l)$$

and

$$\varphi_0 \wedge \exists x_1 \dots \exists x_m (\neg Rx_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} \neg x_k = x_l)$$

are satisfiable.

The class H_0 introduced above is not free. The class of graphs is free, and for any relational τ containing at least one relation symbol of arity ≥ 2 , the class of all τ -structures is a free parametric class.

Proposition 4.3.5 *Let H be a nontrivial free parametric class. Then almost all structures in H are rigid, that is, $u(\text{RIG}|H) = 1$.*

For H the class of all structures we prove the proposition in the appendix of this section.

Together with Theorem 4.3.4 we now obtain the following

Corollary 4.3.6 *Let H be a nontrivial free parametric class. Then the labeled and the unlabeled asymptotic probabilities with respect to H coincide.* \square

We extend this result to:

Proposition 4.3.7 *Let H be any nontrivial parametric class. Then the labeled and the unlabeled asymptotic probabilities with respect to H coincide.*

Before giving a proof we state a consequence. The class H satisfies the *unlabeled 0–1 law* for the set Φ of sentences if for all $\varphi \in \Phi$,

$$u(\varphi|H) = 1 \quad \text{or} \quad u(\varphi|H) = 0.$$

Then we have by 4.2.3:

Theorem 4.3.8 *Let H be a nontrivial parametric class. Then H satisfies the unlabeled 0–1 law for $L_{\infty\omega}^\omega$ and hence, for FO.* \square

Proof (of 4.3.7). By 4.3.6 we only have to consider the nonfree case. First we prove the claim for vocabularies containing only unary relation symbols. So suppose $\tau = \{R_1, \dots, R_m\}$ with unary R_1, \dots, R_m , and let φ_0 be a nontrivial

parametric sentence. For $\alpha : \{1, \dots, m\} \rightarrow \{0, 1\}$ set $R^\alpha x := \varphi_1 \wedge \dots \wedge \varphi_m$, where $\varphi_i = R_i x$ if $\alpha(i) = 1$, and $\varphi_i = \neg R_i x$ if $\alpha(i) = 0$. Then there are distinct $\alpha_1, \dots, \alpha_t$ such that

$$\varphi_0 \text{ and } \forall x (R^{\alpha_1} x \vee \dots \vee R^{\alpha_t} x)$$

are logically equivalent (note that any boolean combination of formulas $R_i x$ can be written as a disjunction of formulas $R^\alpha x$). We assume $t \geq 2$, as the case $t = 1$ is trivial. Since for $i = 1, \dots, t$ and $k \geq 0$ the sentence $\exists^{>k} x R^{\alpha_i} x$ is a consequence of the set $T_{\text{rand}}(\varphi_0)$ of extension axioms compatible with φ_0 , we have

$$(1) \ l(\exists^{=k} x R^{\alpha_i} x \mid \varphi_0) = 0.$$

The isomorphism type of a model \mathcal{A} of φ_0 is determined by (n_1, \dots, n_t) , where $n_i := |\{a \in A \mid \mathcal{A} \models R^{\alpha_i} x[a]\}|$. Using induction on n and – in the induction step – on t , it is easy to show that

$$(2) \text{ the number of } t\text{-tuples } (n_1, \dots, n_t) \text{ such that } n_1 + \dots + n_t = n \text{ equals the binomial coefficient } \binom{n+t-1}{t-1}, \text{ a polynomial in } n \text{ of degree } t-1.$$

$$\text{Hence, } U_n(\varphi_0) = \binom{n+t-1}{t-1}.$$

By 2.3.12 and 3.3.25(b), any sentence of $L_{\infty\omega}^\omega[\tau]$ is equivalent to a finite boolean combination of sentences of the form $\exists^{=k} x R^\alpha x$ (where $k \geq 0$ and $\alpha : \{1, \dots, m\} \rightarrow \{0, 1\}$). Thus it suffices to show that

$$u(\exists^{=k} x R^\alpha x \mid \varphi_0) = l(\exists^{=k} x R^\alpha x \mid \varphi_0) \in \{0, 1\}.$$

If $\alpha \neq \alpha_1, \dots, \alpha \neq \alpha_t$, then $u(\exists^{=k} x R^\alpha x \mid \varphi_0) = l(\exists^{=k} x R^\alpha x \mid \varphi_0) = 0$ or 1 depending on whether $k > 0$ or $k = 0$. Let $\alpha = \alpha_i$. Then by (2), $U_n(\exists^{=k} x R^\alpha x \wedge \varphi_0) = \binom{n-k+t-2}{t-2}$, a polynomial in n of degree $t-2$, and $U_n(\varphi_0) = \binom{n+t-1}{t-1}$, a polynomial in n of degree $t-1$. Hence,

$$u(\exists^{=k} x R^\alpha x \mid \varphi_0) = \lim_{n \rightarrow \infty} \frac{U_n(\exists^{=k} x R^\alpha x \wedge \varphi_0)}{U_n(\varphi_0)} = 0,$$

and thus by (1), $u(\exists^{=k} x R^\alpha x \mid \varphi_0) = l(\exists^{=k} x R^\alpha x \mid \varphi_0) = 0$.

Finally, we turn to arbitrary vocabularies. Let H be a nontrivial parametric class which is not free. By the definition of freeness, if $\mathcal{A} \in H$ then any bijection of the universe that preserves the induced unary relations $\{a \mid R^A a \dots a\}$ for $R \in \tau$, is an automorphism of \mathcal{A} . Hence, the counting arguments in the preceding special case remain valid. \square

Exercise 4.3.9 Let $\sigma = \{P\}$ and $\tau = \{P, R\}$ where P, R are binary. Show for the FO $[\sigma]$ -sentences $\varphi := \exists x \exists y Pxy$ and $\psi := \text{“}P \text{ is an ordering”} \vee \neg \exists x \exists y Pxy$ that $l^\sigma(\varphi|\psi) = 1$, $u^\sigma(\varphi|\psi) = \frac{1}{2}$ and $l^\tau(\varphi|\psi) = u^\tau(\varphi|\psi) = 1$. \square

4.3.1 Appendix

In the following we prove (part of) Proposition 4.3.5. Fix a vocabulary τ and $n \geq 1$, and let $\mathcal{A}, \mathcal{B}, \dots$ range over τ -structures with universe $\{1, \dots, n\}$ and π, ρ, \dots over permutations of $\{1, \dots, n\}$. We set

$$\begin{aligned} \text{Aut}(\mathcal{A}) &:= \{\pi \mid \pi : \mathcal{A} \cong \mathcal{A}\} \\ \text{Str}(\pi) &:= \{\mathcal{A} \mid \pi : \mathcal{A} \cong \mathcal{A}\}. \end{aligned}$$

Lemma 4.3.10 $U_n(\tau) \cdot n! = \sum_{\pi} \|\text{Str}(\pi)\|$.

Proof. $U_n(\tau)$ is the number of equivalence classes of the relation \sim where

$$\mathcal{A} \sim \mathcal{B} \quad \text{iff} \quad \mathcal{A} \cong \mathcal{B}.$$

Clearly,

$$\mathcal{A} \cong \mathcal{B} \quad \text{implies} \quad \|\text{Aut}(\mathcal{A})\| = \|\text{Aut}(\mathcal{B})\|.$$

Given \mathcal{A} and π , once more we let \mathcal{A}_{π} be the structure \mathcal{A} with $\pi : \mathcal{A} \cong \mathcal{A}_{\pi}$. We have already remarked that

$$\mathcal{A}_{\pi} = \mathcal{A}_{\rho} \quad \text{iff} \quad \pi^{-1} \circ \rho \in \text{Aut}(\mathcal{A}).$$

Thus, $\|\{\mathcal{B} \mid \mathcal{A} \cong \mathcal{B}\}\|$ is the index of the subgroup $\text{Aut}(\mathcal{A})$ in the group of all permutations of $\{1, \dots, n\}$; hence,

$$\|\text{Aut}(\mathcal{A})\| = \frac{n!}{\|\{\mathcal{B} \mid \mathcal{A} \cong \mathcal{B}\}\|}.$$

Therefore, for fixed \mathcal{A} ,

$$\sum_{\mathcal{B}, \mathcal{B} \cong \mathcal{A}} \|\text{Aut}(\mathcal{B})\| = \sum_{\mathcal{B}, \mathcal{B} \cong \mathcal{A}} \|\text{Aut}(\mathcal{A})\| = n!.$$

Taking into account that there are $U_n(\tau)$ many equivalence classes with respect to \cong , we have

$$(1) \quad \sum_{\mathcal{B}} \|\text{Aut}(\mathcal{B})\| = U_n(\tau) \cdot n!.$$

On the other hand,

$$(2) \quad \sum_{\mathcal{B}} \|\text{Aut}(\mathcal{B})\| = \|\{(\pi, \mathcal{B}) \mid \pi \in \text{Aut}(\mathcal{B})\}\| = \sum_{\pi} \|\text{Str}(\pi)\|.$$

(1) and (2) yield the desired equation. □

Let M be finite and $f : M \rightarrow M$ be a bijection. Denote by $c(f)$ the number of f -cycles,⁴ by $\text{spt}(f)$ the *support* of f ,

$$\text{spt}(f) := \{a \in M \mid f(a) \neq a\},$$

and set $s(f) := \|\text{spt}(f)\|$. Since $a \in M \setminus \text{spt}(f)$ gives rise to the f -cycle $\{a\}$ and since the f -cycle of any $a \in \text{spt}(f)$ has at least two elements, we have

$$(3) \quad c(f) \leq \|M\| - s(f) + \frac{s(f)}{2} \leq \|M\| - \frac{s(f)}{2}.$$

Proposition 4.3.11 *Let τ be a relational vocabulary that contains at least one relation symbol of arity ≥ 2 . Then $u(\text{RIG}) = 1$.*

Proof. Obviously, we can assume $\|\tau\| = 1$. For simplicity, we restrict ourselves to $\tau = \{E\}$ with binary E . By 4.3.3 it suffices to show that $\lim_{n \rightarrow \infty} \frac{L_n(\tau)}{U_n(\tau) \cdot n!} = 1$, or equivalently, that

$$\lim_{n \rightarrow \infty} \frac{U_n(\tau) \cdot n!}{L_n(\tau)} = 1.$$

Clearly,

$$(4) \quad L_n(\tau) = 2^{n^2}.$$

Fix n and remind our convention that π, ρ, \dots denote permutations of $\{1, \dots, n\}$ and $\mathcal{A}, \mathcal{B}, \dots$ denote τ -structures with universe $\{1, \dots, n\}$.

Each π induces a permutation $\tilde{\pi}$ of $\{1, \dots, n\} \times \{1, \dots, n\}$,

$$\tilde{\pi}((i, j)) := (\pi(i), \pi(j)).$$

If π is an automorphism of \mathcal{A} and $\{\bar{a}, \tilde{\pi}(\bar{a}), \tilde{\pi}(\tilde{\pi}(\bar{a})), \dots\}$ a $\tilde{\pi}$ -cycle, then

$$(5) \quad E^{\mathcal{A}}\bar{a} \text{ iff } E^{\mathcal{A}}\tilde{\pi}(\bar{a}) \text{ iff } E^{\mathcal{A}}\tilde{\pi}(\tilde{\pi}(\bar{a})) \dots$$

that is, we have $E^{\mathcal{A}}\bar{b}$ for all elements \bar{b} of the $\tilde{\pi}$ -cycle of \bar{a} , or for none. Therefore, π is an automorphism for exactly $2^{c(\tilde{\pi})}$ many τ -structures,

$$(6) \quad \|\text{Str}(\pi)\| = 2^{c(\tilde{\pi})}.$$

By the preceding lemma, (4), and (6), we obtain

$$\frac{U_n(\tau) \cdot n!}{L_n(\tau)} = \frac{\sum_{\pi} \|\text{Str}(\pi)\|}{2^{n^2}} = \sum_{\pi} 2^{c(\tilde{\pi}) - n^2}.$$

We must show that $\sum_{\pi} 2^{c(\tilde{\pi}) - n^2} \rightarrow 1$. Since for π the identity on $\{1, \dots, n\}$ we have $c(\tilde{\pi}) = n^2$, this is equivalent to

⁴ For $a \in M$ the set $\{a, f(a), f(f(a)), \dots\}$ is the f -cycle of a .

$$\sum_{\pi \neq id} 2^{c(\tilde{\pi})-n^2} \rightarrow 0.$$

For any π ,

$$\text{spt}(\pi) \times \{1, \dots, n\} \subseteq \text{spt}(\tilde{\pi});$$

hence, by (3), $c(\tilde{\pi}) \leq n^2 - \frac{s(\pi) \cdot n}{2}$ and thus,

$$\sum_{\pi \neq id} 2^{c(\tilde{\pi})-n^2} \leq \sum_{\pi \neq id} 2^{-\frac{s(\pi) \cdot n}{2}}.$$

For $k = 2, \dots, n$ the number of permutations π with $s(\pi) = k$ is $\leq \binom{n}{k} k! \leq n^k$. Therefore, for $n > 2 \cdot \log n$, we have

$$\begin{aligned} \sum_{\pi \neq id} 2^{-\frac{s(\pi) \cdot n}{2}} &\leq \sum_{k=2}^n n^k \cdot 2^{-\frac{k \cdot n}{2}} = \sum_{k=2}^n 2^{-\frac{1}{2}k \cdot (n - 2 \cdot \log n)} \\ &\leq (n-1) \cdot 2^{-(n-2 \cdot \log n)} \end{aligned}$$

(for the last inequality note that $k = 2$ gives the largest summand of the third sum). Since $(n-1) \cdot 2^{-(n-2 \cdot \log n)} \rightarrow 0$ we obtain the desired result. \square

4.4 Examples and Consequences

In the present section we give several examples concerning graphs and draw some consequences for fragments of second-order logic. We start with some general remarks.

Fix a nontrivial parametric sentence φ_0 . By the results of section 4.2, $T_{\text{rand}}(\varphi_0)$ has a uniquely determined countable model $\mathcal{R}(\varphi_0)$, the random model of φ_0 . Some results of the preceding sections are summarized in:

Proposition 4.4.1 *For an $L_{\infty\omega}^\omega$ -sentence φ the following are equivalent:*

- (i) $T_{\text{rand}}(\varphi_0) \models \varphi$
- (ii) $\mathcal{R}(\varphi_0) \models \varphi$
- (iii) $l(\varphi|\varphi_0) = 1$
- (iv) $u(\varphi|\varphi_0) = 1$. \square

Using the extension axioms in $T_{\text{rand}}(\varphi_0)$ we show:

Proposition 4.4.2 (a) *Let \mathcal{B} be a finite model of φ_0 . Then almost all finite models of φ_0 contain a substructure isomorphic to \mathcal{B} .*

(b) *Let \mathcal{B} be a finite model of φ_0 and π_0 an isomorphism of a substructure \mathcal{A} of \mathcal{B} into $\mathcal{R}(\varphi_0)$. Then π_0 can be extended to an isomorphism of \mathcal{B} into $\mathcal{R}(\varphi_0)$.*

Proof. (a) Let \mathcal{B} consist of s elements and let φ_0^s be the conjunction of φ_0 and the finitely many r -extension axioms of $T_{\text{rand}}(\varphi_0)$ with $r \leq s$. Clearly, any model of φ_0^s contains a substructure isomorphic to \mathcal{B} . As $l(\varphi_0^s|\varphi_0) = 1$

the claim follows.

(b) Let $A = \{\bar{a}\}$ and $B = \{\bar{a}, \bar{b}\}$. Then (note that $\varphi_{A, \bar{a}}^0 = \varphi_{B, \bar{a}}^0$)

$$\forall \bar{v}(\varphi_{A, \bar{a}}^0(\bar{v}) \rightarrow \exists \bar{w} \varphi_{B, \bar{a}\bar{b}}^0(\bar{v}, \bar{w}))$$

is a consequence of the extension axioms in $T_{\text{rand}}(\varphi_0)$ and hence, $\mathcal{R}(\varphi_0)$ is a model of it. Since $A \models \varphi_{A, \bar{a}}^0[\bar{a}]$ and $\varphi_{A, \bar{a}}^0$ is quantifier-free, we have $\mathcal{R}(\varphi_0) \models \varphi_{A, \bar{a}}^0[\pi_0(\bar{a})]$, and therefore there are \bar{d} in $\mathcal{R}(\varphi_0)$ such that $\mathcal{R}(\varphi_0) \models \varphi_{B, \bar{a}\bar{b}}^0[\pi_0(\bar{a}), \bar{d}]$. So $\bar{a}\bar{b} \mapsto \pi_0(\bar{a})\bar{d}$ is the required isomorphism. \square

By a back and forth argument we get:

Proposition 4.4.3 *If \mathcal{B} and \mathcal{C} are isomorphic finite substructures of $\mathcal{R}(\varphi_0)$, then there is an automorphism of $\mathcal{R}(\varphi_0)$ mapping \mathcal{B} onto \mathcal{C} .*

Proof. Assume that $\mathcal{R}(\varphi_0)$ has universe $\{a_0, a_1, a_2, \dots\}$ and that $\pi : \mathcal{B} \cong \mathcal{C}$. Take the finite substructure \mathcal{B}'_0 of $\mathcal{R}(\varphi_0)$ with $\mathcal{B}'_0 = \mathcal{B} \cup \{a_0\}$. By part (b) of the preceding proposition, there is an isomorphism $\pi'_0 : \mathcal{B}'_0 \cong \mathcal{C}'_0$ for a suitable $\mathcal{C}'_0 \subseteq \mathcal{R}(\varphi_0)$ that extends π . Similarly, we get $\pi_0 : \mathcal{B}_0 \cong \mathcal{C}_0$ with $\pi'_0 \subseteq \pi_0$, $\mathcal{B}_0 \subseteq \mathcal{R}(\varphi_0)$, $\mathcal{C}_0 \subseteq \mathcal{R}(\varphi_0)$, and $\mathcal{C}_0 = \mathcal{C}'_0 \cup \{a_0\}$. Continuing this way we obtain a sequence $\pi_0 \subseteq \pi_1 \subseteq \pi_2 \subseteq \dots$ with $a_i \in \text{do}(\pi_i) \cap \text{rg}(\pi_i)$. Thus, $\pi := \bigcup_{i \geq 0} \pi_i$ is the desired automorphism of $\mathcal{R}(\varphi_0)$. \square

Graphs

Let φ_G be a parametric sentence axiomatizing the class of graphs. $T_{\text{rand}}(\varphi_G)$ is equivalent to $T_{\text{rand}, G} := \{\varphi_G, \varphi_{\geq 2}\} \cup$

$$\{\forall \text{ distinct } x_1 \dots x_n y_1 \dots y_l \exists z (\bigwedge_{i=1}^n E x_i z \wedge \bigwedge_{i=1}^l (\neg E y_i z \wedge \neg y_i = z)) \mid n + l \geq 1\}.$$

In fact, the sentences in $T_{\text{rand}, G}$ are implied by the extension axioms in $T_{\text{rand}}(\varphi_G)$; hence, any model of $T_{\text{rand}}(\varphi_G)$ is a model of $T_{\text{rand}, G}$. On the other hand, a back and forth argument using the axioms in $T_{\text{rand}, G}$ shows that any two of its models and hence, any model of $T_{\text{rand}, G}$ and any model of $T_{\text{rand}}(\varphi_G)$ are partially isomorphic; therefore, any model of $T_{\text{rand}, G}$ is a model of $T_{\text{rand}}(\varphi_G)$.

It is well known that a graph cannot be planar if it contains the subgraph \mathcal{K}_5 , a clique with 5 elements. Thus, by part (a) of 4.4.2, we get:

Proposition 4.4.4 *Almost all finite graphs are not planar.* \square

Moreover, we have:

Proposition 4.4.5 *$\mathcal{R}(\varphi_G)$, the random graph, and almost all finite graphs \mathcal{G} are connected, the diameter $D(\mathcal{G}) := \max\{d(a, b) \mid a, b \in \mathcal{G}\}$ being 2.*

Proof. Note that, by the extension axioms, $\mathcal{R}(\varphi_G)$ and almost all finite graphs are models of

$$(*) \quad \exists x \exists y \neg Exy \wedge \forall x \forall y \exists z (Exz \wedge Eyz).$$

But any graph \mathcal{G} satisfying this sentence is connected with $D(\mathcal{G}) = 2$. \square

Even though “connectedness” is not expressible in first-order logic, the first-order sentence $(*)$ of the preceding proof is a property of almost all graphs that implies “connectedness”. The situation for “rigidity” is different:

Proposition 4.4.6 (a) *Almost all finite graphs are rigid.*

(b) $\mathcal{R}(\varphi_G)$ *is not rigid.*

(c) *No $L_{\infty\omega}^\omega$ -definable property of almost all graphs implies rigidity.*

Proof. For part (a) cf. 4.3.5 and 4.3.2. Part (b) is an immediate consequence of 4.4.3 (for \mathcal{B} and \mathcal{C} take two substructures of cardinality one). Part (c) follows from (b) because any $L_{\infty\omega}^\omega$ -definable property of almost all graphs holds in $\mathcal{R}(\varphi_G)$ (cf. 4.4.1). \square

Fragments of Second-Order Logic

Connectedness of graphs can be expressed by a Π_1^1 -sentence, for example by

$$\varphi_{\text{CONN}} := \forall X (\forall x Xx \vee \forall x \neg Xx \vee \exists x \exists y (Xx \wedge \neg Xy \wedge Exy)).$$

Nonrigidity is expressible by

$$\begin{aligned} & \exists X \forall x \forall y \forall u \forall v \exists z_1 \exists z_2 \exists w (Xz_1x \wedge Xxz_2 \wedge \neg Xw \wedge \\ & ((Xxy \wedge Xuv) \rightarrow ((x = u \leftrightarrow y = v) \wedge (Exu \leftrightarrow Eyv))))), \end{aligned}$$

a so-called $\Sigma_1^1(\forall^*\exists^*)$ -sentence, that is, a sentence of the form

$$\exists X_1 \dots \exists X_s \forall y_1 \dots \forall y_m \exists z_1 \dots \exists z_l \chi$$

with $s, m, l \in \mathbb{N}$ and χ quantifier-free. Similarly, $\Sigma_1^1(\exists^*\forall^*)$ -sentences have the form

$$\exists X_1 \dots \exists X_s \exists y_1 \dots \exists y_m \forall z_1 \dots \forall z_l \chi$$

where χ is quantifier-free.

Part (a) of the following proposition generalizes the fact mentioned above that connectedness is implied by a first-order property of almost all graphs. Part (b) shows that nonrigidity cannot be expressed by a $\Sigma_1^1(\exists^*\forall^*)$ -sentence (otherwise almost all finite graphs would be nonrigid).

Proposition 4.4.7 *Suppose that φ_0 is nontrivial parametric.*

(a) *Let φ be a Π_1^1 -sentence. If $\mathcal{R}(\varphi_0) \models \varphi$ then there is a first-order sentence ψ such that*

$$u(\psi|\varphi_0) = 1 \quad \text{and} \quad \models \psi \rightarrow \varphi.$$

(b) Let φ be a $\Sigma_1^1(\exists^*\forall^*)$ -sentence. If $\mathcal{R}(\varphi_0) \models \varphi$ then there is a first-order sentence ψ such that

$$u(\psi|\varphi_0) = 1 \quad \text{and} \quad \models_{\text{fin}} \psi \rightarrow \varphi.$$

Proof. (a) Assume that $\mathcal{R}(\varphi_0) \models \varphi$ holds for the Π_1^1 -sentence

$$\varphi = \forall X_1 \dots \forall X_s \chi,$$

where χ contains no second-order quantifiers. Then the set $T_{\text{rand}}(\varphi_0) \cup \{\neg\chi\}$ of $\tau \cup \{X_1, \dots, X_s\}$ -sentences has no model (otherwise, by the Löwenheim-Skolem theorem, it would have a countable model whose τ -reduct would be (isomorphic to) the unique countable model $\mathcal{R}(\varphi_0)$ of $T_{\text{rand}}(\varphi_0)$; but then $\mathcal{R}(\varphi_0) \models \exists X_1 \dots \exists X_s \neg\chi$, contrary to $\mathcal{R}(\varphi_0) \models \forall X_1 \dots \forall X_s \chi$). By compactness, there is a finite subset T_0 of $T_{\text{rand}}(\varphi_0)$ such that $T_0 \cup \{\neg\chi\}$ is not satisfiable. Let ψ be the conjunction of the sentences in T_0 . Then $u(\psi|\varphi_0) = 1$ and $\models \psi \rightarrow \chi$, hence $\models \psi \rightarrow \forall X_1 \dots \forall X_s \chi$.

(b) Suppose that for the $\Sigma_1^1(\exists^*\forall^*)$ -sentence

$$\varphi = \exists X_1 \dots \exists X_s \exists \bar{x} \forall \bar{y} \chi$$

with quantifier-free χ we have $\mathcal{R}(\varphi_0) \models \varphi$, say, $(\mathcal{R}(\varphi_0), X_1, \dots, X_s) \models \exists \bar{x} \forall \bar{y} \chi$. Choose \bar{a} in $\mathcal{R}(\varphi_0)$ such that

$$(*) \quad (\mathcal{R}(\varphi_0), X_1, \dots, X_s) \models \forall \bar{y} \chi[\bar{a}]$$

and denote by \mathcal{A}_0 the submodel of $\mathcal{R}(\varphi_0)$ with universe $\{\bar{a}\}$. Since $\exists \bar{x} \varphi_{\bar{a}}^0(\bar{x})$ holds in $\mathcal{R}(\varphi_0)$, there is a ψ that is the conjunction of φ_0 and of finitely many extension axioms compatible with φ_0 such that $\models \psi \rightarrow \exists \bar{x} \varphi_{\bar{a}}^0(\bar{x})$. Obviously, $u(\psi|\varphi_0) = 1$. We show that $\models_{\text{fin}} \psi \rightarrow \varphi$.

So let \mathcal{B} be a finite model of ψ . Choose \bar{d} in \mathcal{B} such that $\mathcal{B} \models \varphi_{\bar{a}}^0[\bar{d}]$. Then $\bar{d} \mapsto \bar{a}$ is an isomorphism from the substructure of \mathcal{B} with universe $\{\bar{d}\}$ into $\mathcal{R}(\varphi_0)$. By 4.4.2(b), there is an extension π of $\bar{d} \mapsto \bar{a}$ that is an isomorphism of \mathcal{B} onto a substructure \mathcal{B}' of $\mathcal{R}(\varphi_0)$. It suffices to show that \mathcal{B}' is a model of φ . By (*), as $\forall \bar{y} \chi$ is universal,

$$(\mathcal{B}', X_1 \cap \mathcal{B}', \dots, X_s \cap \mathcal{B}') \models \forall \bar{y} \chi[\bar{a}],$$

hence $(\mathcal{B}', X_1 \cap \mathcal{B}', \dots, X_s \cap \mathcal{B}') \models \exists \bar{x} \forall \bar{y} \chi$, thus $\mathcal{B}' \models \exists X_1 \dots \exists X_s \exists \bar{x} \forall \bar{y} \chi$. \square

As a consequence we get:

Theorem 4.4.8 *Let φ_0 be nontrivial parametric and let φ be a $\Sigma_1^1(\exists^*\forall^*)$ -sentence.*

- (a) *If $\mathcal{R}(\varphi_0) \models \varphi$ then $u(\varphi|\varphi_0) = 1$.*
- (b) *If $\mathcal{R}(\varphi_0) \not\models \varphi$ then $u(\varphi|\varphi_0) = 0$.*

Proof. (a) If $\mathcal{R}(\varphi_0) \models \varphi$ then, by part (b) of the proposition, there is a first-order sentence ψ such that $u(\psi|\varphi_0) = 1$ and $\models_{\text{fin}} \psi \rightarrow \varphi$. In particular, $u(\varphi|\varphi_0) = 1$.

For (b), assume that $\mathcal{R}(\varphi_0) \models \neg\varphi$. Since $\neg\varphi$ is (logically equivalent to) a Π_1^1 -sentence there is, by part (a) of the proposition, a first-order sentence ψ such that $u(\psi|\varphi_0) = 1$ and $\models \psi \rightarrow \neg\varphi$. Therefore, $u(\neg\varphi|\varphi_0) = 1$ and hence, $u(\varphi|\varphi_0) = 0$. \square

By 4.3.7, the preceding results remain valid for the labeled probability as well. Hence, we have:

Corollary 4.4.9 $\Sigma_1^1(\exists^*\forall^*)$ satisfies the labeled and the unlabeled 0–1 law with respect to nontrivial parametric classes. \square

The satisfiability problem for $\exists^*\forall^*$ -sentences is decidable, and we have just seen that $\Sigma_1^1(\exists^*\forall^*)$ has the 0–1 law. This is a special case of a general phenomenon: It has turned out that the satisfiability problem for a prefix class Φ of first-order logic is decidable just in case the 0–1 law holds for $\Sigma_1^1\Phi$ ($:= \{\exists\bar{R}\varphi \mid \varphi \in \Phi\}$), cf. [126].

Orderings

In the preceding exposition we dealt with parametric classes; the following examples show, among other things, that the class of orderings is not parametric.

Examples 4.4.10 (a) Let $\tau = \{<, P\}$ with unary P and let $\mathcal{O} = \mathcal{O}[\tau]$ be the class of τ -structures ordered by $<$. Denote by φ a first-order sentence expressing that the first element of the ordering is in P , say, $\varphi = \exists x(Px \wedge \forall y \neg y < x)$. Then

$$u(\varphi|\mathcal{O}) = l(\varphi|\mathcal{O}) = \frac{1}{2}$$

(argue as in 4.1.1(a)). By 4.2.3, \mathcal{O} is not parametric.

(b) For $\tau = \{<\}$ consider the class ORD of orderings. By 2.3.6, for any first-order sentence φ we have $u(\varphi|\text{ORD}) = l(\varphi|\text{ORD}) = 1$ iff $(\{0, \dots, 2^k\}, <) \models \varphi$, where k is the quantifier rank of φ . Therefore, ORD satisfies the labeled and the unlabeled 0–1 law for FO. On the other hand, ORD does not satisfy the (labeled or unlabeled) 0–1 law for $L_{\infty\omega}^2$. In fact, the probabilities $l(\varphi|\text{ORD})$ and $u(\varphi|\text{ORD})$ do not exist for any $L_{\infty\omega}^2$ -sentence φ expressing that the ordering has an even number of elements. (As φ one can take $\bigvee_{n \geq 1} \chi_{2n}$, where χ_{2n} are the FO²-formulas introduced in Example 3.3.1(a).) \square

4.5 Probabilities of Monadic Second Order Properties

Already in 4.1.1(c) we saw that, in general, the (labeled or unlabeled) asymptotic probabilities do not always exist for second-order sentences. In the

present section we show that this is true even for monadic Σ_1^1 -sentences, that is, for sentences of the form

$$\exists X_1 \dots \exists X_m \psi$$

where X_1, \dots, X_m are unary and ψ is first-order. Moreover, we shall show the corresponding result for nontrivial free parametric classes.

Theorem 4.5.1 *Let φ_0 be a nontrivial free parametric sentence in a relational vocabulary τ_0 . Then there is a monadic Σ_1^1 -sentence φ of vocabulary τ_0 such that the labeled asymptotic probability $l(\varphi|\varphi_0)$ does not exist.*

Since the class of all finite τ -structures is a nontrivial parametric class in case τ contains an at least binary relation symbol, we obtain

Corollary 4.5.2 *Let τ be a vocabulary that contains an at least binary relation symbol. Then there is a monadic Σ_1^1 -sentence of vocabulary τ without labeled asymptotic probability. \square*

Proof (of 4.5.1). Let φ_0 be a nontrivial free parametric sentence and H its class of finite models. Then the following lemma 4.5.3 tells us that there is a first-order formula $\kappa(x, y, X, Y, Z, U)$ with unary relation variables X, Y, Z, U such that

$$\exists X \exists Y \exists Z \exists U (\text{“}\kappa(-, -, X, Y, Z, U) \text{ is an ordering of the universe”})$$

almost surely holds on H . Therefore, with a new unary relation variable V , the sentence

$$\exists X \exists Y \exists Z \exists U \exists V (\text{“}\kappa(-, -, X, Y, Z, U) \text{ is an ordering of the universe whose first element belongs to } V \text{ and whose last element does not belong to } V \text{”} \wedge \forall x \forall y (\text{“if } y \text{ is the } \kappa\text{-successor of } x \text{ then } (Vx \leftrightarrow \neg Vy)\text{”}))$$

almost surely holds for the structures in H of even cardinality and almost surely gets wrong for the structures in H of odd cardinality. So it has no labeled asymptotic probability (recall that φ_0 is nontrivial and thus, has models in all cardinalities). \square

Lemma 4.5.3 *There is a first-order formula $\kappa(x, y, X, Y, Z, U)$ with unary relation variables X, Y, Z, U such that with*

$$\varphi := \exists X \exists Y \exists Z \exists U (\text{“}\kappa(-, -, X, Y, Z, U) \text{ is an ordering of the universe”})$$

the labeled asymptotic probability $l(\varphi|\varphi_0)$ equals 1.

Proof. As φ_0 is free, there is a relation symbol R in τ_0 of arity $r \geq 2$, an $m \geq 2$, and a surjection $i: \{1, \dots, r\} \rightarrow \{1, \dots, m\}$ such that both

$$\varphi_0 \wedge \exists x_1 \dots \exists x_m (R x_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} x_k \neq x_l)$$

and

$$\varphi_0 \wedge \exists x_1 \dots \exists x_m (\neg R x_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} x_k \neq x_l)$$

are satisfiable. We set

$$\psi(x_1, \dots, x_m) := R x_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} x_k \neq x_l$$

and

$$\psi^\neg(x_1, \dots, x_m) := \neg R x_{i(1)} \dots x_{i(r)} \wedge \bigwedge_{1 \leq k < l \leq m} x_k \neq x_l.$$

Then, among the 0-isomorphism types of m distinct elements satisfiable in some model of φ_0 , there are types implying $\psi(x_1, \dots, x_m)$ and types implying $\psi^\neg(x_1, \dots, x_m)$. Hence, the fraction p of types implying $\psi(x_1, \dots, x_m)$ satisfies $0 < p < 1$. Interchanging ψ and ψ^\neg , if necessary, we may assume that

$$0 < p \leq \frac{1}{2}.$$

Intuitively, if we take distinct elements a_1, \dots, a_m in a model \mathcal{A} of φ_0 , the probability for $\psi[a_1, \dots, a_m]$ to hold in \mathcal{A} is p .

We first consider the case $m = 2$ and write $\psi(x, y)$ for $\psi(x_1, x_2)$. The idea behind our argumentation is as follows: We show that with high probability models of φ_0 contain a subset of logarithmic size ordered by $\psi(x, y)$. The ordering can suitably be extended to an ordering of the whole universe; this will be done in two steps.

We come to the details. Given n , we set

$$\begin{aligned} c &:= -\log_p n, \\ r &:= \lfloor \sqrt{\frac{1}{2}c} \rfloor, \text{ the integer part of } \sqrt{\frac{1}{2}c}, \\ s &:= \lfloor 3c \rfloor, \\ t &:= \lfloor \frac{n}{2\sqrt{c}} \rfloor - 1, \end{aligned}$$

$$X := \{1, \dots, r\} \quad \text{and} \quad Z := \{n - s + 1, \dots, n\}.$$

Note that for sufficiently large n , the sets X and Z are disjoint.

In the following we equip the set of models of φ_0 over $\{1, \dots, n\}$ with the uniform probability distribution.

Claim 1. Almost surely in models of φ_0 over $\{1, \dots, n\}$ there is a set Y of power r disjoint from $X \cup Z$ such that

$$\psi^Y_{\leq}(x, y) := Yx \wedge Yy \wedge \exists z (Xz \wedge \psi(z, x) \wedge \neg\psi(z, y))$$

defines an ordering on Y .

Claim 2. Almost surely in models of φ_0 over $\{1, \dots, n\}$, for any subset Y of $\{1, \dots, n\}$ of cardinality r which is disjoint from $X \cup Z$, we have

- (1) $\forall x \forall y (Zx \wedge Zy \wedge x \neq y \rightarrow \exists z (Yz \wedge \neg(\psi(z, x) \leftrightarrow \psi(z, y))))$;
- (2) with $U := (Y \cup Z)^c$:

$$\forall x \forall y ((Ux \wedge Uy \wedge x \neq y) \rightarrow \exists z (Zz \wedge \neg(\psi(z, x) \leftrightarrow \psi(z, y)))).$$

Then we are done: If $\psi_{<}^Y(x, y)$ defines an ordering on Y and (1), (2) of Claim 2 are valid, then

$$\psi_{<}^Z(x, y) := x \neq y \wedge Zx \wedge Zy \wedge \text{“for the } \psi_{<}^Y\text{-smallest element } z \text{ of } Y \text{ with } \neg(\psi(z, x) \leftrightarrow \psi(z, y)) \text{ we have } \psi(z, x)\text{”}$$

defines an ordering on Z and

$$\psi_{<}^U(x, y) := x \neq y \wedge Ux \wedge Uy \wedge \text{“for the } \psi_{<}^Z\text{-smallest element } z \text{ of } Z \text{ with } \neg(\psi(z, x) \leftrightarrow \psi(z, y)) \text{ we have } \psi(z, x)\text{”}$$

defines an ordering on U .

Therefore, with

$$\begin{aligned} \kappa(x, y, X, Y, Z, U) := & \text{“} X, Y, Z \text{ are disjoint and } U = (Y \cup Z)^c \text{”} \wedge \\ & (\psi_{<}^Y(x, y) \vee \psi_{<}^Z(x, y) \vee \psi_{<}^U(x, y)) \\ & \vee (Yx \wedge \neg Yy) \vee (Zx \wedge Uy), \end{aligned}$$

claims 1 and 2 yield that

$$l(\exists X \exists Y \exists Z \exists U \text{“} \kappa(-, -, X, Y, Z, U) \text{ defines a linear ordering”} \mid \varphi_0) = 1.$$

So, we have φ as claimed in the statement of the lemma.

We give the proofs of Claim 1 and Claim 2. To prove Claim 1, for $i \in \{1, \dots, t\}$ we set

$$Y(i) := \{r \cdot i + 1, \dots, r \cdot i + r\}.$$

Note that the $Y(i)$ are pairwise disjoint and disjoint from both X and Z (as $r \cdot t + r < \frac{n}{2} < n - s + 1$). We show that almost surely in models of φ_0 over $\{1, \dots, n\}$ there is an $i \in \{1, \dots, t\}$ such that

$$(*) \quad \text{for all } a, b \in X, \psi(a, r \cdot i + b) \text{ holds iff } a \leq b.$$

Then almost surely there is a Y , namely $Y = Y(i)$ for an i satisfying $(*)$, such that for all $d_1, d_2 \in Y$

there is an $a \in X$ such that $\psi(a, d_1)$ holds and $\neg\psi(a, d_2)$ holds iff $d_2 < d_1$,

that is, $\psi^Y_{\leq}(x, y)$ defines an ordering on Y .

For $i \in \{1, \dots, n\}$ let q be the probability that i satisfies $(*)$ (where we refer to the uniform probability distribution on the set of models of φ_0 over $\{1, \dots, n\}$). Then

$$q \geq p^{r^2} \geq p^{\frac{1}{2}c} = \frac{1}{\sqrt{n}}.$$

Hence, the probability that no $i \in \{1, \dots, t\}$ satisfies $(*)$ is

$$\begin{aligned} &\leq \left(1 - \frac{1}{\sqrt{n}}\right)^t \leq \left(1 - \frac{1}{\sqrt{n}}\right)^{\frac{t}{\sqrt{2c}}-2} \\ &\leq \left(\left(1 - \frac{1}{\sqrt{n}}\right)^{\sqrt{n}}\right)^{\frac{\sqrt{n}}{2\sqrt{c}} - \frac{2}{\sqrt{n}}} \rightarrow_{n \rightarrow \infty} 0. \end{aligned}$$

We now come to a proof of Claim 2. Given Y , the probability of the failure of (1) is

$$\leq \binom{s}{2} \cdot (p^2 + (1-p)^2)^r \leq 5 \cdot c^2 \cdot (p^2 + (1-p)^2)^{\sqrt{\frac{1}{2}c}-1} \rightarrow_{n \rightarrow \infty} 0,$$

and the probability for the failure of (2) is

$$\begin{aligned} &\leq \binom{n}{2} \cdot (p^2 + (1-p)^2)^s \\ &\leq n^2 \cdot (p^2 + (1-p)^2)^{\log_p n^{-3}-1} \\ &\leq n^2 \cdot (p^2 + (1-p)^2)^{\log_{p^2+(1-p)^2} n^{-3}-1} \quad (\text{note that } p^2 + (1-p)^2 \geq p) \\ &\leq \frac{1}{n \cdot (p^2 + (1-p)^2)} \rightarrow_{n \rightarrow \infty} 0. \end{aligned}$$

Altogether the case $m = 2$ is settled.

We finally sketch a proof for the case $m > 2$. The model construction method for parametric classes given in section 4.2 shows that for distinct elements a_1, \dots, a_m the probability for $\psi(a_1, \dots, a_m)$ to hold in a model over $\{1, \dots, n\}$ remains the same if we fix a_3, \dots, a_m . Therefore, the preceding proof for the case $m = 2$ remains valid with the following changes:

- (i) We consider structures over $\{1, \dots, n, n+1, \dots, n+m-2\}$.
- (ii) We fix x_3 by $n+1, \dots, x_m$ by $n+m-2$.
- (iii) We replace $U = (Y \cup Z)^c$ by the complement of $Y \cup Z$ within $\{1, \dots, n\}$.
- (iv) We put $n+1, \dots, n+m-2$ at the end of the ordering defined on $\{1, \dots, n\}$ for the case $m = 2$.

This leads to the following changes in the definition of φ :

- (i) We add “ $\exists x_3 \dots \exists x_m$ ” to the prefix.
- (ii) We add the conjunct “ $\bigwedge_{3 \leq k < l \leq m} x_k \neq x_l$ ” to the kernel.
- (iii) We define $\kappa(x, y, x_3, \dots, x_m, X, Y, Z, U)$ as in case $m = 2$ with x_3, \dots, x_m as parameters stemming from $\psi(x, y, x_3, \dots, x_m)$ and then replace it by the disjunction of the following formulas

$$\begin{aligned} & \bigwedge_{3 \leq k \leq m} x \neq x_k \wedge \bigvee_{3 \leq k \leq m} y = x_k, \\ & \bigvee_{3 \leq k < m} \bigvee_{k < j \leq m} (x = x_k \wedge y = y_j), \\ & \bigwedge_{3 \leq k \leq m} (x \neq x_k \wedge y \neq x_k) \wedge \kappa(x, y, x_3, \dots, x_m, X', Y', Z', U') \end{aligned}$$

where, e.g., X' is $X \setminus \{x_3, \dots, x_k\}$.

□

Exercise 4.5.4 Let τ contain the relation symbol $<$ and at least one further relation symbol of arity ≥ 2 . Let $\mathcal{O}[\tau]$ be the class of ordered τ -structures. Then there is a first-order φ such that $l(\varphi|\mathcal{O}[\tau])$ does not exist. (Hint: Argue as in the preceding proof, but now let intervals take over the role of the subsets X, Y, Z and U .) □

Notes 4.5.5 The 0–1 law for first-order logic was independently proven by Glebskij, Logan, Liogonkij, Talanov [47] and Fagin [37], its extension to parametric classes is due to Oberschelp [123], and its extension to $L_{\infty\omega}^\omega$ to Kolaitis and Vardi [109]. The papers [107, 108] contain, among other things, the results on fragments of second-order logic presented in section 4.4. The last section is based on [98] (cf. also [99, 141]) and the Appendix 4.3.1 on [38]. Survey articles on 0–1 laws are [21, 72, 111, 146].

5. Satisfiability in the Finite

A classical question having its origin in the decision problem for first-order logic, asks for specific classes Φ of sentences whether they have the finite model property, i.e., whether every satisfiable sentence of Φ has a finite model. We are going to consider two examples where we can demonstrate the methodological usefulness of techniques and results we have developed so far.

A first-order sentence φ expressing that $<$ is a partial ordering without maximal elements is satisfiable but has no finite model. As φ we can take (1) or (2):

$$(1) \quad \forall x \neg x < x \wedge \forall x \forall y \forall z ((x < y \wedge y < z) \rightarrow x < z) \wedge \forall x \exists y x < y$$

$$(2) \quad \forall x \forall y \forall z \exists u (\neg x < x \wedge ((x < y \wedge y < z) \rightarrow x < z) \wedge x < u).$$

The sentence in (1) uses only three variables, the sentence in (2) is a $\forall^3\exists$ -sentence. These sentences are “best” possible ones: We show in section 5.1 that every satisfiable sentence with at most two variables has a finite model and in section 5.2 that the same holds for every satisfiable $\forall^2\exists^*$ -sentence without the equality sign.¹

5.1 Finite Model Property of FO^2

We prove that every satisfiable sentence with at most two variables in a relational vocabulary has a finite model. As a consequence, every sentence of modal logic has a finite model (cf. 5.1.8). We remark that one can remove the restriction on constants (cf. 5.1.7), but the result is not valid for “vocabularies” with function symbols: Consider the sentence with two variables

$$\forall x \forall y (f(x) = f(y) \rightarrow x = y) \wedge \exists y \forall x \neg f(x) = y$$

expressing that f is injective but not surjective.

Fix a relational vocabulary τ and let $x := v_1$ and $y := v_2$. For the purpose of this section we say that a first-order formula (possibly containing second-order variables) is *normal*, if it has the form

¹ Furthermore, any satisfiable universal sentence has a finite model.

$$\forall x \forall y \psi \wedge \bigwedge_{i=1}^r \forall x \exists y \psi_i$$

where $\psi, \psi_i \in \text{FO}^2$ are quantifier-free. We shall use the next lemma to restrict ourselves to normal sentences when proving the finite satisfiability of satisfiable FO^2 -sentences.

Lemma 5.1.1 *Every sentence $\exists \bar{X}(\varphi \wedge \forall x \forall y \psi)$, where φ is normal and $\psi \in \text{FO}^2$, is equivalent to a sentence of the form $\exists \bar{Y} \chi$, where χ is normal.*

Proof. We proceed by induction on the number of quantifiers in ψ . If ψ contains no quantifiers, the result is immediate. In the induction step we show how to eliminate a quantifier in ψ . So let, say, $\forall x \psi_0$ be a subformula of ψ with quantifier-free ψ_0 . Then, ψ is logically equivalent to

$$\exists X (\forall y (Xy \leftrightarrow \forall x \psi_0) \wedge \forall x \forall y \psi')$$

where ψ' results from ψ by replacing $\forall x \psi_0$ by Xy , and hence, it is logically equivalent to

$$\exists X (\forall x \forall y (Xy \rightarrow \psi_0) \wedge \forall y \exists x (\neg \psi_0 \vee Xy) \wedge \forall x \forall y \psi')$$

and thus to

$$\exists X (\forall x \forall y (Xy \rightarrow \psi_0) \wedge \forall x \exists y (\neg \psi_0 \binom{yx}{xy} \vee Xx) \wedge \forall x \forall y \psi')$$

$(\psi_0 \binom{yx}{xy})$ is obtained from ψ_0 by simultaneously replacing *all* occurrences of x and y by y and x , respectively).

Altogether, $\exists \bar{X}(\varphi \wedge \forall x \forall y \psi)$ is equivalent to

$$\exists \bar{X} \exists X (\varphi \wedge \forall x \forall y (Xy \rightarrow \psi_0) \wedge \forall x \exists y (\neg \psi_0 \binom{yx}{xy} \vee Xx) \wedge \forall x \forall y \psi')$$

where the first conjunct is normal and ψ' has less quantifiers than ψ . By induction hypothesis, we obtain our claim. \square

Corollary 5.1.2 *Every sentence of FO^2 is equivalent to a sentence of the form $\exists \bar{Y} \chi$, where χ is normal.*

Proof. Given an FO^2 -sentence ψ , apply the preceding lemma to $\exists \bar{X}(\varphi \wedge \forall x \forall y \psi)$, where \bar{X} is the empty sequence and $\varphi := \forall x \forall y x = x$. \square

We now formulate the main result of this section.

Theorem 5.1.3 *Every satisfiable first-order sentence with at most two variables in a relational vocabulary has a finite model.*

Proof. Let φ be such a sentence. We apply the corollary above. Since $\exists\bar{Y}\chi$ and χ have models over the same universes, we may assume that φ has the form

$$\varphi = \forall x \forall y \psi \wedge \bigwedge_{i=1}^r \forall x \exists y \psi_i$$

where $\psi, \psi_i \in \text{FO}^2$ are quantifier-free. Moreover, we may suppose that

$$\text{for } i = 1, \dots, r, \quad \psi_i \models x \neq y$$

since over structures with at least two elements $\forall x \exists y \psi_i(x, y)$ is equivalent to

$$\forall x \exists y (x \neq y \wedge (\psi_i(x, y) \vee \psi_i(x, x))).$$

Let \mathcal{A} be a model of φ . An element $a \in A$ is a *king* (in \mathcal{A}) if there is no other element b of A with the same 0-isomorphism type, i.e., $\varphi_{\mathcal{A},a}^0 = \varphi_{\mathcal{A},b}^0$. If $\mathcal{A} \models \psi_i[a, b]$, we call b a *child* of a (in \mathcal{A}), more exactly, an *i -child*. For $a \in A$ and $i \in \{1, \dots, r\}$ we let a^i be a fixed i -child of a . Then, $a \neq a^i$. We set

$$C := \bigcup_{a \in A, a \text{ king}} \{a, a^1, \dots, a^r\}.$$

Clearly, C is finite. We show that there is a \mathcal{B} such that

- (i) $B = C \cup (\{\varphi_{\mathcal{A},a}^0 \mid a \in A \text{ no king}\} \times \{1, \dots, r\} \times \{0, 1, 2\})$.
- (ii) Each 0-isomorphism type of a pair of elements of B is realized in \mathcal{A} .
- (iii) For $i = 1, \dots, r$, all elements of B have an i -child in \mathcal{B} .

Then \mathcal{B} is a model of $\forall x \forall y \psi$ by (ii), and of $\bigwedge_{i=1}^r \forall x \exists y \psi_i$ by (iii); thus by (i), \mathcal{B} is a finite model of φ .

To define \mathcal{B} , we fix the 0-isomorphism type of all pairs of elements of B in a suitable way to ensure (ii) and (iii). The rest – in case τ contains relation symbols of arity ≥ 3 – can be fixed in an arbitrary way.

Step 1. For $a, b \in C, a \neq b$, we set

$$\varphi_{\mathcal{B},a,b}^0 := \varphi_{\mathcal{A},a,b}^0.$$

Step 2. Let $b \in B$. We aim at providing children for b in \mathcal{B} . So let $i \in \{1, \dots, r\}$.

- If $b \in C$ and b is a king or b has an i -child in \mathcal{A} that lies in C , b has an i -child in \mathcal{B} because of Step 1.
- If $b = a^j$ for a king a , but b has no i -child in C , we let $b' := (\varphi_{\mathcal{A},(a^j)^i}^0, i, 0)$ be an i -child of b in \mathcal{B} by setting

$$\varphi_{\mathcal{B},b,b'}^0 := \varphi_{\mathcal{A},a^j,(a^j)^i}^0.$$

(In case there are several possibilities for a and j , we fix one choice; and we also do so in similar situations.)

- If $b = (\varphi_{\mathcal{A},a}^0, j, k)$ (and, hence, $a \in A$ is not a king in \mathcal{A}) and a^i is a king in \mathcal{A} , we let a^i be an i -child of b in \mathcal{B} by setting

$$\varphi_{\mathcal{B},b,a^i}^0 := \varphi_{\mathcal{A},a,a^i}^0.$$

- If $b = (\varphi_{\mathcal{A},a}^0, j, k)$ and if a^i is not a king in \mathcal{A} , we let $b' := (\varphi_{\mathcal{A},a^i}^0, i, (k+1) \pmod{3})$ be an i -child of b in \mathcal{B} by setting

$$\varphi_{\mathcal{B},b,b'}^0 := \varphi_{\mathcal{A},a,a^i}^0.$$

In all of these cases, by fixing a type $\varphi_{\mathcal{B},a,b}^0$, we of course at the same time also fix $\varphi_{\mathcal{B},b,a}^0$.

Step 3. If, e.g., for $d \in C$, $b := (\varphi_{\mathcal{A},a}^0, j, k)$, and $b' := (\varphi_{\mathcal{A},a'}^0, j', k')$ the 0-isomorphism type of (d, b) or of (b, b') has not been fixed in the first two steps, we set

$$\varphi_{\mathcal{B},d,b}^0 := \varphi_{\mathcal{A},d,a}^0 \quad \text{or} \quad \varphi_{\mathcal{B},b,b'}^0 := \varphi_{\mathcal{A},a,a'}^0,$$

respectively. Note that the definitions we have made do not contradict each other, as for $c \in C$ we have $\varphi_{\mathcal{B},c}^0 = \varphi_{\mathcal{A},c}^0$ and for $b = (\varphi_{\mathcal{A},a}^0, j, k)$ we have $\varphi_{\mathcal{B},b}^0 = \varphi_{\mathcal{A},a}^0$. By construction, (ii) and (iii) are satisfied. \square

Corollary 5.1.4 *For any relational vocabulary τ , the set Φ of logically valid first-order sentences with at most two variables is decidable.*

Proof. By the completeness theorem for first-order logic the set Φ is enumerable. For its “complement”

$$\Phi^{\text{nv}} := \{\varphi \text{ FO}^2[\tau]\text{-sentence} \mid \varphi \text{ is not logically valid}\}$$

we have by the preceding theorem that

$$\Phi^{\text{nv}} = \{\varphi \text{ FO}^2[\tau]\text{-sentence} \mid \neg\varphi \text{ has a finite model}\}.$$

Therefore, Φ^{nv} is enumerable too, and hence, Φ is decidable. \square

We remarked after Corollary 4.4.9 that the classical prefix classes with a decidable satisfaction problem correspond to fragments of Σ_1^1 with a 0–1 law. Even though we just have shown that FO^2 has a decidable satisfaction problem, the 0–1 law does not hold for the class $\Sigma_1^1\text{-FO}^2 := \{\exists \bar{X}\varphi \mid \varphi \in \text{FO}^2\}$, cf. [114].

An analysis of the proof of the main theorem leads to the statement of the following exercise.

Exercise 5.1.5 Show that to any first-order sentence φ in at most two variables one can effectively assign a natural number m such that:

- φ has a model if and only if φ has a model of power $\leq m$.
- If φ has a model of power $\geq m$ then φ has a model in each cardinality $\geq m$. □

Exercise 5.1.6 Let $\tau = \{<\}$ with $<$ binary. Show that for all $n \geq 1$ there is a satisfiable $\text{FO}^2[\tau]$ -sentence having only models of cardinality $\geq n$. (Hint: Use the formulas introduced in part (a) of 3.3.1, but note that the class of orderings is not axiomatizable in FO^2 .) □

Exercise 5.1.7 Show that any satisfiable sentence with at most two variables in a vocabulary possibly containing constants has a finite model. (Hint: Let φ be such a sentence and denote by C the set of constants. For any n -ary relation symbol P and any map $\pi : \{1, \dots, n\} \rightarrow C \cup \{x, y\}$ introduce a new 2-ary relation symbol P^π . Apply the theorem to the sentence obtained from φ by replacing any subformula $P\pi(1) \dots \pi(n)$ by $P^\pi xy$ and treating equations and quantifiers appropriately.) □

Exercise 5.1.8 Show that any satisfiable sentence of modal logic has a finite model. The sentences of modal logic are built up from propositional variables p_1, p_2, \dots by means of the propositional connectives, say \neg and \vee , and the “necessity operator” \Box . A *structure* (or *frame*) \mathcal{M} is a triple (A, R, \mathcal{F}) where A is a nonempty set, the set of “states”, R is a binary relation on A , and $\mathcal{F} : \{p_1, p_2, \dots\} \rightarrow \text{Pow}(A)$, $\text{Pow}(A)$ being the power set of A . For $s \in A$, the satisfaction relation $\mathcal{M} \models_s \alpha$ is defined by induction on the complexity of α :

$$\begin{array}{lll}
 \mathcal{M} \models_s p_i & \text{iff} & s \in \mathcal{F}(p_i) \\
 \mathcal{M} \models_s \neg \alpha & \text{iff} & \text{not } \mathcal{M} \models_s \alpha \\
 \mathcal{M} \models_s (\alpha \vee \beta) & \text{iff} & \mathcal{M} \models_s \alpha \text{ or } \mathcal{M} \models_s \beta \\
 \mathcal{M} \models_s \Box \alpha & \text{iff} & \mathcal{M} \models_t \alpha \text{ for all } t \text{ such that } Rst.
 \end{array}$$

If $\mathcal{M} \models_s \alpha$ for all $s \in A$, then \mathcal{M} is said to be a *model* of α . To solve the exercise, let the propositional variables occurring in α be among p_1, \dots, p_n . Associate with α a formula $\varphi_\alpha(x) \in \text{FO}^2[\{R, P_1, \dots, P_n\}]$, where P_1, \dots, P_n are unary, such that for all (A, R, \mathcal{F}) and $s \in A$,

$$(A, R, \mathcal{F}) \models_s \alpha \quad \text{iff} \quad (A, R, \mathcal{F}(p_1), \dots, \mathcal{F}(p_n)) \models \varphi_\alpha[s]. \quad \square$$

5.2 Finite Model Property of $\forall^2\exists^*$ -Sentences

We fix a relational vocabulary τ . In the preceding section we have proved that every FO^2 -sentence has models in the same cardinalities as a sentence of the form

$$\forall x \forall y \psi \wedge \bigwedge_{i=1}^r \forall x \exists y \psi_i$$

with quantifier-free ψ, ψ_i . This sentence is equivalent to

$$(*) \quad \forall x \forall y \exists y_1 \dots \exists y_r (\psi(x, y) \wedge \bigwedge_{i=1}^r \psi_i(x, y_i)).$$

We then have proved the finite model property for these sentences.

By a $\forall^2\exists^*$ -sentence we mean a first-order sentence of the form

$$\forall x_1 \forall x_2 \exists y_1 \dots \exists y_k \psi'$$

where $k \geq 0$ and ψ' is quantifier-free. In particular, the sentence $(*)$ is a $\forall^2\exists^*$ -sentence. In this section we extend the result about the finite model property for sentences of the form $(*)$ to arbitrary $\forall^2\exists^*$ -sentences, however under the proviso that they have models without kings. (Recall that $a \in A$ is a king in the structure \mathcal{A} if for no $b \in A, b \neq a$, we have $\varphi_a^0 = \varphi_b^0$.)

Theorem 5.2.1 *Suppose that τ is a relational vocabulary. If ψ is a $\forall^2\exists^*$ -sentence which has a model without kings, then it has a finite model.*

We first draw a consequence. For a structure \mathcal{A} in a relational vocabulary τ and $l \geq 2$, denote by $\mathcal{A} \times l$ the structure which for every element of A contains l duplicates; more precisely: $\mathcal{A} \times l$ is the τ -structure with universe $A \times \{0, \dots, l-1\}$ such that for any n -ary R in τ ,

$$R^{\mathcal{A} \times l} := \{((a_1, i_1), \dots, (a_n, i_n)) \mid R^{\mathcal{A}} a_1 \dots a_n, 0 \leq i_1, \dots, i_n \leq l-1\}.$$

Clearly, $\mathcal{A} \times l$ is a structure without kings, and a routine proof by induction shows that

$$\mathcal{A} \models \psi \quad \text{iff} \quad \mathcal{A} \times l \models \psi$$

holds for all sentences ψ without (the) equality (sign). As a corollary of the above theorem we therefore obtain:

Corollary 5.2.2 *Suppose that τ is a relational vocabulary and ψ is a $\forall^2\exists^*$ -sentence without equality. If ψ is satisfiable then it has a finite model.*

As in 5.1.4 for FO^2 we now get:

Corollary 5.2.3 *The set of logically valid $\forall^2\exists^*$ -sentences without equality in a relational vocabulary is decidable. \square*

Proof (of 5.2.1). In models with at least two elements a $\forall^2\exists^*$ -sentence

$$\forall v_1 \forall v_2 \exists v_3 \dots \exists v_k \psi'(v_1, \dots, v_k)$$

is equivalent to the sentence

$$\forall v_1 \forall v_2 \exists x_3 \dots \exists x_k \exists z_3 \dots \exists z_k (\neg v_1 = v_2 \rightarrow (\psi'(v_1, v_1, x_3, \dots, x_k) \wedge \psi'(v_1, v_2, z_3, \dots, z_k))),$$

and, for example, a sentence of the form

$$\forall v_1 \forall v_2 \exists x \exists y \psi'(v_1, v_2, x, y)$$

is equivalent to

$$\forall v_1 \forall v_2 \exists x \exists y ((\psi'(v_1, v_2, x, x) \vee \psi'(v_1, v_2, x, y)) \wedge \neg x = y).$$

Therefore, we can assume that our $\forall^2\exists^*$ -sentence ψ has the form

$$\forall v_1 \forall v_2 \exists v_3 \dots \exists v_k (\neg v_1 = v_2 \rightarrow (\psi'(v_1, \dots, v_k) \wedge \bigwedge_{3 \leq i < j \leq k} \neg v_i = v_j))$$

where ψ' is quantifier-free.

Choose a model \mathcal{A} of ψ without kings and let

$$S := \{\varphi_a^0 \mid a \in A\} \quad \text{and} \quad T := \{\varphi_{ab}^0 \mid a, b \in A, a \neq b\}$$

be the 0-isomorphism types of elements and of pairs of elements of \mathcal{A} , respectively. If $\rho(v_1, \dots, v_l)$ is a 0-isomorphism type of any l -tuple and $1 \leq m, n \leq l$ with $m \neq n$, let $\rho_m(v_1)$ and $\rho_{m,n}(v_1, v_2)$ be the induced 0-isomorphism types of v_m and of v_m, v_n , respectively; in particular, for any \mathcal{B} and $b_1, \dots, b_l \in B$,

$$\mathcal{B} \models \rho[b_1, \dots, b_l] \text{ implies } \varphi_{b_m}^0 = \rho_m \text{ and } \varphi_{b_m b_n}^0 = \rho_{m,n}.$$

As \mathcal{A} has no kings, we get:

- (1) For all $\varphi, \varphi' \in S$ there is a $\chi \in T$ such that $\varphi = \chi_1$ and $\varphi' = \chi_2$.

Moreover, since \mathcal{A} is a model of ψ , we have:

- (2) For every $\chi \in T$ there is a 0-isomorphism type $\rho(v_1, \dots, v_k)$ with
- (a) $\rho_i \in S$ for $i = 1, \dots, k$
 - (b) $\rho_{m,n} \in T$ for $1 \leq m < n \leq k$ and $\rho_{1,2} = \chi$
 - (c) $\models \rho \rightarrow \psi'(v_1, \dots, v_k)$.

To get the statement of the theorem it suffices to show:

- (+) Suppose that S and T are nonempty sets of 0-isomorphism types of elements and of pairs of elements, respectively, satisfying (1) and (2). Then ψ has a finite model.

Let $s := \|S\|$ and $t := \|T\|$ and fix an ordering on S . To prove (+), we give a method to construct, for every $n \geq k$, structures \mathcal{B} with universe $\{1, 2, \dots, n \cdot s\}$ (afterwards we show that with nonvanishing probability these structures are models of ψ). The 0-isomorphism types of elements are fixed by a deterministic algorithm while the 0-isomorphism types of tuples of more than one element are chosen randomly. The exact construction procedure of \mathcal{B} reads as follows:

- (i) If $a \in \{1, 2, \dots, n \cdot s\}$ and $a = i \cdot s + j$ for some i, j such that $0 \leq i < n$ and $1 \leq j \leq s$, ensure that φ_a^0 is equal to the j -th element in S .
- (ii) If $a, b \in \{1, 2, \dots, n \cdot s\}$, $1 \leq a < b \leq n \cdot s$, choose at random a χ in $\{\chi \in T \mid \chi_1 = \varphi_a^0, \chi_2 = \varphi_b^0\}$ (this set is nonempty by (1)) and ensure that $\varphi_{ab}^0 = \chi$.
- (iii) If R is an m -ary relation symbol in τ , define the truth value of $Ra_1 \dots a_m$ at random for any $a_1, \dots, a_m \in \{1, 2, \dots, n \cdot s\}$ containing at least three and at most k distinct members.
- (iv) If R is an m -ary relation symbol in τ , define the truth value of $Ra_1 \dots a_m$ to be “false” if $a_1 \dots a_m$ contains more than k distinct members.

Let $\text{Str}(n)$ be the collection of possible values of \mathcal{B} with $\{1, 2, \dots, n \cdot s\}$ as universe. Equip $\text{Str}(n)$ with the uniform probability distribution μ . Let $\bar{a} = a_1 \dots a_k$ denote pairwise distinct elements of $\{1, 2, \dots, n \cdot s\}$ and let d be the number of formulas $Rv_{i_1} \dots v_{i_m}$ where $R \in \tau$ and $\{v_{i_1}, \dots, v_{i_m}\}$ contains at least three and at most k distinct variables.

Claim 1. Suppose that $\chi \in T$ and that the 0-isomorphism type $\rho(v_1, \dots, v_k)$ satisfies (2) with respect to χ . Then the conditional probability

$$\mu(\varphi_{\bar{a}}^0 = \rho \mid \varphi_{a_1 a_2}^0 = \chi, \varphi_{a_i}^0 = \rho_i \text{ for } i = 3, \dots, k)$$

is $\geq \delta$, where $\delta = (\frac{1}{i})^{\binom{k}{2}-1} \cdot (\frac{1}{2})^d$, that is,

$$\mu(\{\mathcal{B} \mid \mathcal{B} \models \rho[\bar{a}]\} \mid \{\mathcal{B} \mid \mathcal{B} \models \chi[a_1, a_2], \mathcal{B} \models \rho_i[a_i] \text{ for } i = 3, \dots, k\}) \geq \delta.$$

Proof. The proof is immediate by the definition of d and the fact that, once $\varphi_{a_i}^0$ for $i = 1, \dots, k$ and $\varphi_{a_1 a_2}^0$ are fixed, 0-isomorphism types according to (ii) have to be chosen for $\binom{k}{2} - 1$ pairs of distinct elements.

Claim 2. Fix a_1, a_2 . Then

$$\mu(\{\mathcal{B} \mid \mathcal{B} \not\models \exists v_3 \dots \exists v_k (\psi'(a_1, a_2, v_3, \dots, v_k) \wedge \bigwedge_{1 \leq i < j \leq k} v_i \neq v_j)\}) \leq (1 - \delta)^f,$$

where f is the integer part of $\frac{n-2}{k-2}$.

Proof. Let $\chi \in T$ and choose (for χ) a corresponding $\rho(v_1, \dots, v_k)$ according to (2). It suffices to prove that the conditional probability

$$(3) \quad \mu(\{\mathcal{B} \mid \mathcal{B} \not\models \exists v_3 \dots \exists v_k \rho(a_1, a_2, v_3, \dots, v_k)\} \mid \varphi_{a_1 a_2}^0 = \chi) \leq (1 - \delta)^f.$$

By (i), in any \mathcal{B} every 0-isomorphism type in S is realized by n ($\geq 2 + f \cdot (k-2)$) distinct elements. Therefore, for $i = 3, \dots, k$ and $j = 1, \dots, f$ there are pairwise distinct elements $a_i^j \in \{1, 2, \dots, n \cdot s\} \setminus \{a_1, a_2\}$ with $\varphi_{a_i^j}^0 = \rho_i$. Under the given conditions the events $\mathcal{B} \models \rho(a_1, a_2, a_3^j, \dots, a_k^j)$ for $1 \leq j \leq f$ are independent (compare the construction procedure). Now,

$$\begin{aligned} & \{\mathcal{B} \mid \mathcal{B} \not\models \exists v_3 \dots \exists v_k \rho(a_1, a_2, v_3, \dots, v_k)\} \\ & \subseteq \{\mathcal{B} \mid \text{for } j = 1, \dots, f, \mathcal{B} \not\models \rho(a_1, a_2, a_3^j, \dots, a_k^j)\}. \end{aligned}$$

Therefore, by Claim 1 we obtain (3).

Since $\{\mathcal{B} \mid \mathcal{B} \not\models \psi\} =$

$$\bigcup_{\substack{a_1, a_2 \\ a_1 \neq a_2}} \{\mathcal{B} \mid \mathcal{B} \not\models \exists v_3 \dots \exists v_k (\psi'(a_1, a_2, v_3, \dots, v_k) \wedge \bigwedge_{1 \leq i < j \leq k} v_i \neq v_j)\},$$

we obtain by Claim 2 that

$$\mu(\{\mathcal{B} \mid \mathcal{B} \not\models \psi\}) \leq n \cdot s \cdot (n \cdot s - 1) \cdot (1 - \delta)^f.$$

As f is the integer part of $\frac{n-2}{k-2}$, for n big enough we have $n \cdot s \cdot (n \cdot s - 1) \cdot (1 - \delta)^f < 1$. Then the probability that \mathcal{B} satisfies ψ is positive, and therefore some member of $\text{Str}(n)$ satisfies ψ . \square

Exercise 5.2.4 Let τ be a vocabulary (which perhaps contains constants). Show that any satisfiable $\forall^2\exists^*$ -sentence without equality has a finite model (argue as in 5.1.7). Conclude that any satisfiable $\exists^*\forall^2\exists^*$ -sentence without equality has a finite model. \square

Notes 5.2.5 Theorem 5.1.3 is due to Mortimer [121], Corollary 5.1.4 for sentences without equality was already known before (cf. [131]). The proof of 5.1.3 given here incorporates simplifications from [55]. Together with [44] it shows that the satisfiability problem for FO^2 is NEXPTIME-complete. The papers [59, 60, 127, 144] discuss the decidability of extensions and fragments of FO^2 . Corollary 5.2.2 goes back to Gödel [48]; the proof given for 5.2.1 is taken from [75]. For formulas with equality the corresponding satisfiability problem is undecidable ([49]; see also [50]). For thorough information on the decidability and the finite model property of fragments of first-order logic we refer the reader to [13].

6. Finite Automata and Logic: A Microcosm of Finite Model Theory

One of the major aims of finite model theory consists in characterizing the queries in a given complexity class by means of a logic in which they can be described. In this way one obtains a new measure of complexity: the complexity of formal descriptions. Moreover, the characterizations allow to translate problems on complexity classes into purely modeltheoretic problems, thus opening them to modeltheoretic methods. In general, this relationship between complexity classes and logics turns out to be useful for both, complexity theory and logic.

Our considerations will show that first-order logic which has played a central role so far is not adequate for descriptive complexity. What is missing? Which features of computations cannot be expressed in first-order logic and which natural extensions of first-order logic capture these features? Once these questions are clear they show the way to stronger systems such as logics with fixed-point operators that we shall treat in the next chapters. These extensions have natural modeltheoretic properties and thus may turn out to be useful in other areas of model theory, too.

The present chapter deals with the – historically first – logical characterization of a complexity class: the characterization of the class of languages recognized by automata by means of monadic second-order logic. It turns out that this simple case already reflects some of the crucial aspects encountered later, at the same time witnessing that the interplay between model theory and complexity theory is fruitful for both sides.

Even though we include all definitions and proofs, a certain familiarity with automata will be helpful.

6.1 Languages Accepted by Automata

We recall some definitions and notations. Let \mathbb{A} be a finite alphabet and \mathbb{A}^* and \mathbb{A}^+ the set of strings (or words) and the set of nonempty strings over \mathbb{A} , respectively. Thus, $\mathbb{A}^* = \mathbb{A}^+ \cup \{\lambda\}$, where λ is the empty word. In automata

theory subsets of \mathbb{A}^+ are called *languages*.¹ A *nondeterministic automaton* M , in short, an NDA (over the alphabet \mathbb{A}) is given by a tuple

$$M = (S, q_0, \delta, F),$$

where S is a finite set, the set of *states*, $q_0 \in S$ is the *initial state*, $F \subseteq S$ is the set of (*accepting* or) *final states* and $\delta \subseteq S \times \mathbb{A} \times S$ is the *transition relation* (intuitively, $(q, a, p) \in \delta$ means: if M is in state q and reads a , then M can pass into p). This relation induces a function $\tilde{\delta} : S \times \mathbb{A}^* \rightarrow \text{Pow}(S)$, where $\text{Pow}(S)$ denotes the power set of S ; $\tilde{\delta}$ is given by

$$\begin{aligned} \tilde{\delta}(q, \lambda) &:= \{q\} \\ \tilde{\delta}(q, wa) &:= \{p \mid (r, a, p) \in \delta \text{ for some } r \in \tilde{\delta}(q, w)\}. \end{aligned}$$

In particular, $\tilde{\delta}(q, a) = \{p \mid (q, a, p) \in \delta\}$ for $a \in \mathbb{A}$. If $\tilde{\delta}(q, a)$ is a singleton for every $a \in \mathbb{A}$, then M is said to be a *deterministic automaton* or, by short, an *automaton*. Clearly, in this case $\tilde{\delta}(q, w)$ is a singleton for any $w \in \mathbb{A}^*$. If $\tilde{\delta}(q, w) = \{p\}$, we simply write $\tilde{\delta}(q, w) = p$ (and similarly, $\delta(q, a) = p$ stands for $\tilde{\delta}(q, a) = \{p\}$).

The language *recognized* (or *accepted*) by the NDA M is defined by

$$L(M) := \{w \in \mathbb{A}^+ \mid \tilde{\delta}(q_0, w) \cap F \neq \emptyset\}.$$

Hence, in case M is deterministic,

$$L(M) = \{w \in \mathbb{A}^+ \mid \tilde{\delta}(q_0, w) \in F\}.$$

Let us state the main result of this chapter, even though we still have not introduced all concepts appearing in it. With respect to the logical characterizations we have in mind, we are mainly interested in the equivalence of (ii) and (vi). The other equivalences are not only useful for the proof, but also interesting in their own.

For a language $L \subseteq \mathbb{A}^+$ the following are equivalent:

- (i) *L is the union of equivalence classes of an invariant equivalence relation on \mathbb{A}^+ of finite index.*
- (ii) *L is recognized by an automaton.*
- (iii) *L is recognized by an NDA.*
- (iv) *L is regular.*
- (v) *L is definable in monadic second-order logic by a Σ_1^1 -sentence.*
- (vi) *L is definable in monadic second-order logic.*

¹ In the literature languages, in particular languages recognized by automata, may contain the empty word. Since to any word we will assign a structure whose cardinality coincides with the length of the word, it is more convenient for our purposes to exclude the empty string.

Note that (ii) \Rightarrow (iii) and (v) \Rightarrow (vi) are trivial. We prove the other implications by a series of propositions starting with those which mainly deal with automata; monadic second-order logic will be treated in the next section.

An equivalence relation \sim on \mathbb{A}^+ is called *invariant* if

$$u, v, w \in \mathbb{A}^+ \text{ and } u \sim v \text{ imply } uw \sim vw.$$

Denote by $[u]$ the equivalence class of u and by \mathbb{A}^+ / \sim the set of equivalence classes. The *index* of \sim is the cardinality of \mathbb{A}^+ / \sim .

Proposition 6.1.1 *Let \sim be an invariant equivalence relation on \mathbb{A}^+ of finite index. Suppose that the language $L \subseteq \mathbb{A}^+$ is the union of equivalence classes,*

$$L = [u_1] \cup \dots \cup [u_r]$$

for some $u_1, \dots, u_r \in \mathbb{A}^+$. Then L is recognized by an automaton.

Proof. Add $[\lambda]$, “the equivalence class of λ ”, as a new object to \mathbb{A}^+ / \sim and define the automaton $M = (S, q_0, \delta, F)$ by

$$S := \mathbb{A}^+ / \sim \cup \{[\lambda]\}, \quad q_0 := [\lambda], \quad \delta([u], a) := [ua], \quad F := \{[u_1], \dots, [u_r]\}.$$

By invariance of \sim , the transition function δ is well-defined. For $u, v \in \mathbb{A}^*$, a simple induction on the length of v shows that $\tilde{\delta}([u], v) = [uv]$, in particular, $\tilde{\delta}([\lambda], v) = [v]$. Therefore,

$$L(M) = \{v \in \mathbb{A}^+ \mid \tilde{\delta}(q_0, v) \in F\} = \{v \in \mathbb{A}^+ \mid [v] \in F\} = [u_1] \cup \dots \cup [u_r] = L. \quad \square$$

For later applications we state the following result:

Lemma 6.1.2 (Pumping Lemma) *Let \sim be an invariant equivalence relation on \mathbb{A}^+ of finite index. Then there is an $n \geq 0$ such that for any word $u \in \mathbb{A}^+$ with $|u| \geq n$ ² there exist $v, w \in \mathbb{A}^+$ and $x \in \mathbb{A}^*$ with*

$$u = vwx, \quad |vw| \leq n, \quad \text{and} \quad vw^k \sim vw \text{ for all } k \geq 0.$$

Hence, by invariance, $vw^k y \sim vw y$ for all $k \geq 0$ and $y \in \mathbb{A}^$.*

Proof. Let l be the index of \sim and set $n := l + 1$. Suppose that $u \in \mathbb{A}^+$, $u = a_1 \dots a_s$, where $a_1, \dots, a_s \in \mathbb{A}$ and $s \geq n$. Then, for some i and j with $1 \leq i < j \leq n$, we have $a_1 \dots a_i \sim a_1 \dots a_j$. Let $v = a_1 \dots a_i$ and $w = a_{i+1} \dots a_j$. Thus $v \sim vw$, and by invariance of \sim , $vw \sim vw^2 \sim vw^3 \sim \dots$. \square

We now turn to a description of the languages recognized by automata in terms of simple, so-called regular expressions. First some notations: The *concatenation* of languages L_1 and L_2 , denoted by $L_1 L_2$, is the set $\{uv \mid u$

² $|u|$ denotes the length of u .

in L_1 and v is in L_2 }. Define $L^1 := L$ and $L^n := L^{n-1}L$ for $n > 1$. The *plus* (or *positive*) *closure* L^+ of L is the set $L^+ := \bigcup_{n \geq 1} L^n$. The *regular expressions* (over \mathbb{A}) are strings over the alphabet $\{\emptyset\} \cup \{\mathbf{a} \mid a \in \mathbb{A}\} \cup \{\cup, +, \cdot, \{\cdot\}$. Together with the languages they denote, the regular expressions are defined recursively as follows:

- (a) \emptyset is a regular expression and denotes the empty set.
- (b) \mathbf{a} is a regular expression and denotes the set $\{a\}$.
- (c) If r and s are regular expressions denoting the languages R and S , respectively, then $(r \cup s)$, (rs) , and r^+ are regular expressions that denote the sets $R \cup S$, RS , and R^+ , respectively.

A language is *regular* if it is denoted by some regular expression. For convenience we often omit parentheses in regular expressions (such as in $r_1 \cup \dots \cup r_k$) that have no influence on the language they denote. Moreover, we assume that $+$ has higher precedence than concatenation or \cup , and that concatenation has higher precedence than \cup .

Proposition 6.1.3 *If L is recognized by an NDA then L is regular.*

Proof. Let L be recognized by the NDA $M = (S, q_0, \delta, F)$ with $S = \{q_0, \dots, q_n\}$. Let L_k^{ij} be the set of all nonempty strings that M can read starting in q_i and ending in q_j without going through any state numbered $\geq k$,

$$L_k^{ij} := \{b_1 \dots b_s \mid s \geq 1, b_1, \dots, b_s \in \mathbb{A}, \text{ there are } i_0, \dots, i_s \text{ such that } i_1, \dots, i_{s-1} < k, i_0 = i, i_s = j, \text{ and } (q_{i_m}, b_{m+1}, q_{i_{m+1}}) \in \delta \text{ for } m < s\}.$$

Since $L(M) = \bigcup_{q_i \in F} L_{n+1}^{0j}$, it suffices to show that all L_k^{ij} are regular. We proceed by induction on k . As $L_0^{ij} = \{a \in \mathbb{A} \mid (q_i, a, q_j) \in \delta\}$ is a subset of \mathbb{A} , say $L_0^{ij} = \{a_1, \dots, a_r\}$, L_0^{ij} is denoted by $(\mathbf{a}_1 \cup \dots \cup \mathbf{a}_r)$, or by \emptyset in case $r = 0$.

For the induction step note that a nonempty string is in L_{k+1}^{ij} , if it can be read without visiting any state numbered $\geq k + 1$, thereby starting in q_i , ending in q_j , and passing through q_k zero times or one or more than one time. Hence,

$$L_{k+1}^{ij} = L_k^{ij} \cup L_k^{ik} L_k^{kj} \cup L_k^{ik} (L_k^{kk})^+ L_k^{kj}.$$

By induction hypothesis, for all i', j' there is a regular expression $r_k^{i'j'}$ denoting $L_k^{i'j'}$. Therefore, L_{k+1}^{ij} is denoted by

$$r_k^{ij} \cup r_k^{ik} r_k^{kj} \cup r_k^{ik} (r_k^{kk})^+ r_k^{kj}. \quad \square$$

6.2 Word Models

We turn to languages definable in monadic second-order logic. In order to bridge the gap between automata and logic, we introduce a correspondence

between words and structures. Once more we fix an alphabet \mathbb{A} and let $\tau(\mathbb{A})$ be the vocabulary $\{<\} \cup \{P_a \mid a \in \mathbb{A}\}$, where $<$ is binary and the P_a are unary. For a given $u \in \mathbb{A}^+$, say $u = a_1 \dots a_n$, we consider structures of the form

$$(B, <, (P_a)_{a \in \mathbb{A}})$$

where the cardinality of B equals the length of u , $<$ is an ordering of B , and P_a corresponds to the positions in u carrying an a ,

$$P_a := \{b \in B \mid \text{for some } j, b \text{ is the } j\text{-th element of } < \text{ and } a_j = a\}.$$

We call them *word models* for u and denote the class of word models for u by K_u . For example, if $\mathbb{A} = \{a, b\}$ and $u = abbab$, the structure

$$(\{1, \dots, 5\}, <, P_a, P_b),$$

where $<$ is the natural ordering on $\{1, \dots, 5\}$, $P_a = \{1, 4\}$, and $P_b = \{2, 3, 5\}$ is a word model for u .

Any two word models for u are isomorphic. For simplicity we therefore often speak of *the* word model for u , denoting it by \mathcal{B}_u .

Note that for $u, v \in \mathbb{A}^+$ we obtain a word model for uv by forming the ordered sum (cf. 1.A3) $\mathcal{B}_u \triangleleft \mathcal{B}_v$.

A language $L \subseteq \mathbb{A}^+$ is *definable in monadic second-order logic*, if there is a sentence φ in $\text{MSO}[\tau(\mathbb{A})]$ such that $\text{Mod}(\varphi) = \bigcup_{u \in L} K_u$, or, more succinctly (but not fully correct), $\text{Mod}(\varphi) = \{\mathcal{B}_u \mid u \in L\}$. The notion of definability in FO is introduced similarly.

Let φ_W be the first-order sentence

$$\varphi_W := \text{“} < \text{ is a total ordering”} \wedge \forall x \bigvee_{a \in \mathbb{A}} P_a x \wedge \bigwedge_{\substack{a, b \in \mathbb{A} \\ a \neq b}} \forall x \neg (P_a x \wedge P_b x).$$

Then, $\text{Mod}(\varphi_W)$ is the class of all word models, $\text{Mod}(\varphi_W) = \{\mathcal{B}_u \mid u \in \mathbb{A}^+\}$, and therefore the language \mathbb{A}^+ is definable in first-order logic.

As the next step in the proof of the main result stated at the beginning of this chapter we shall show that any regular language is definable in monadic second-order logic. First we introduce some notations: Let $\psi_{\min}(x)$ and $\psi_{\max}(x)$ be first-order formulas defining the first and the last element of the ordering, respectively:

$$\psi_{\min}(x) := \forall y \neg y < x, \quad \psi_{\max}(x) := \forall y \neg x < y.$$

For any formula φ of MSO and variables x and y let $\varphi^{[x,y]}$ be a formula expressing that the closed interval $[x, y]$ satisfies φ , and similarly, $\varphi^{]x,y]}$ a formula expressing that the half-open interval $]x, y]$ satisfies φ . Such formulas can be obtained from φ by relativizing the first-order quantifiers to the interval, the main clause of an inductive definition being (for $z \neq x, z \neq y$)

$$\begin{aligned}
(\exists z\varphi)^{[x,y]} &:= \exists z(x \leq z \wedge z \leq y \wedge \varphi^{[x,y]}) \\
(\exists z\varphi)^{]x,y]} &:= \exists z(x < z \wedge z \leq y \wedge \varphi^{]x,y]}).
\end{aligned}$$

Proposition 6.2.1 *Any regular language is definable in monadic second-order logic by a Σ_1^1 -sentence.*

Proof. First we prove by induction on the length of the regular expression r that there is a sentence φ_r of MSO defining the language denoted by r . Afterwards we show that we can replace φ_r by a Σ_1^1 -sentence.

$$\begin{aligned}
\varphi_\emptyset &:= \exists x \neg x = x, & \varphi_a &:= \varphi_W \wedge \exists x \forall y (y = x \wedge P_a x) \\
\varphi_{(r \cup s)} &:= \varphi_W \wedge (\varphi_r \vee \varphi_s) \\
\varphi_{(rs)} &:= \varphi_W \wedge \text{“the universe is partitioned into two intervals satisfying } \varphi_r \text{ and } \varphi_s, \text{ respectively”} \\
&= \varphi_W \wedge \exists x \exists y \exists z (\psi_{\min}(x) \wedge y < z \wedge \psi_{\max}(z) \wedge \varphi_r^{[x,y]} \wedge \varphi_s^{]y,z]}) \\
\varphi_{r^+} &:= \varphi_W \wedge \text{“there is a set of right endpoints of intervals, which partition the universe, all parts satisfying } \varphi_r \text{”} \\
&= \varphi_W \wedge \exists X (\exists y (Xy \wedge \psi_{\max}(y)) \wedge \\
&\quad \exists x \exists y (\psi_{\min}(x) \wedge Xy \wedge \forall z (z < y \rightarrow \neg Xz) \wedge \varphi_r^{[x,y]}) \wedge \\
&\quad \forall x \forall y \\
&\quad ((x < y \wedge Xx \wedge Xy \wedge \forall z (x < z < y \rightarrow \neg Xz)) \rightarrow \varphi_r^{]x,y]})).
\end{aligned}$$

Now, one obtains a Σ_1^1 -sentence by inductively bringing all existential second-order quantifiers to the front. In general, a monadic second-order formula $\forall \bar{x} \exists Y \chi$ with first-order χ is not equivalent to a monadic Σ_1^1 -formula (cf. 2.4.8). However, in the case of the formula in the last two lines of φ_{r^+} we can argue as follows: Suppose that φ_r is equivalent to $\exists Y_1 \dots \exists Y_m \chi$. In models of φ_W (and only these are relevant) the formula in the last two lines is equivalent to

$$\exists Y_1 \dots \exists Y_m \forall x \forall y ((x < y \wedge Xx \wedge Xy \wedge \forall z (x < z < y \rightarrow \neg Xz)) \rightarrow \chi^{]x,y]})$$

(for the nontrivial implication of the equivalence piece Y_1, \dots, Y_m together from corresponding subsets chosen in the (disjoint) intervals). \square

With the next proposition we close the last gap in the proof of the main result.

Proposition 6.2.2 *Let $L \subseteq \mathbb{A}^+$ be definable in monadic second-order logic. Then there is an invariant equivalence relation on \mathbb{A}^+ of finite index such that L is a union of equivalence classes.*

Proof. Assume that $\text{Mod}(\varphi) = \{B_u \mid u \in L\}$ where φ is a sentence of MSO. Let m be the quantifier rank of φ and define \sim on \mathbb{A}^+ by

$$u \sim v \quad \text{iff} \quad B_u \equiv_m^{\text{MSO}} B_v$$

(recall that $\mathcal{A} \equiv_m^{\text{MSO}} \mathcal{B}$ means that \mathcal{A} and \mathcal{B} satisfy the same sentences of MSO of quantifier rank $\leq m$). Clearly, \sim is an equivalence relation. Since, up to logical equivalence, there are only finitely many sentences of quantifier rank $\leq m$ (cf. 3.1.3), the relation \sim is of finite index. By definition of m ,

$$\mathcal{B}_u \models \varphi \text{ and } u \sim v \text{ imply } \mathcal{B}_v \models \varphi.$$

Thus, $L = \bigcup\{[u] \mid u \in \mathbb{A}^+, \mathcal{B}_u \models \varphi\}$. Finally, \sim is invariant: Assume $u \sim v$ and $w \in \mathbb{A}^+$. Then $\mathcal{B}_u \equiv_m^{\text{MSO}} \mathcal{B}_v$. Since \equiv_m^{MSO} is preserved by ordered sums (cf. 3.1.4), we obtain

$$\mathcal{B}_{uw} \cong \mathcal{B}_u \triangleleft \mathcal{B}_w \equiv_m^{\text{MSO}} \mathcal{B}_v \triangleleft \mathcal{B}_w \cong \mathcal{B}_{vw}$$

that is, $uw \sim vw$. □

Summing up we have shown the result stated at the beginning of section 6.1:

Theorem 6.2.3 *For a language $L \subseteq \mathbb{A}^+$ the following are equivalent:*

- (i) L is the union of equivalence classes of an invariant equivalence relation on \mathbb{A}^+ of finite index.
- (ii) L is recognized by an automaton.
- (iii) L is recognized by an NDA.
- (iv) L is regular.
- (v) L is definable in monadic second-order logic by a Σ_1^1 -sentence.
- (vi) L is definable in monadic second-order logic. □

In particular, this theorem shows that a language is accepted by an automaton just in case it is definable in monadic second-order logic. Another description of languages accepted by automata is given by means of the regular expressions. Is that a logical description too? What makes a logic? We shall address this problem later.

6.3 Examples and Applications

By the following examples and applications we try to exhibit the methodological usefulness of Theorem 6.2.3.

Proposition 6.3.1 (a) *The class of languages over \mathbb{A} accepted by automata is closed under the boolean operations (complementation and union).*

- (b) (Pumping Lemma). *Let L be accepted by an automaton. Then there is $n \geq 0$ such that for any $u \in \mathbb{A}^+$ with $|u| \geq n$ there exist $v, w \in \mathbb{A}^+$ and $x \in \mathbb{A}^*$ with*

$$u = vwx, |vw| \leq n, \text{ and for } k \geq 0 \text{ and } y \in \mathbb{A}^* : vw^k y \in L \text{ iff } vwy \in L.$$

Proof. Part (a) holds, since monadic second-order logic is closed under the boolean connectives \neg and \vee , and part (b) is a reformulation of the Pumping Lemma 6.1.2. Both results also have simple direct proofs. \square

Example 6.3.2 Let $\mathbb{A} = \{a\}$. Identify $\underbrace{a \dots a}_{n\text{-times}}$ with the natural number n and thus, \mathbb{A}^+ with the set \mathbb{N}_+ of positive natural numbers. A subset L of \mathbb{N}_+ is accepted by an automaton iff L is *ultimately periodic*, that is, if there are $p, r \in \mathbb{N}_+$ such that for all $m \geq p$,

$$m + r \in L \quad \text{iff} \quad m \in L.$$

In fact, assume first that L is accepted by an automaton. By the Pumping Lemma there are $n, j, r \in \mathbb{N}_+$ and $l \geq 0$ with $n = j + r + l$ such that for all $k \geq 0$ and $s \in \mathbb{N}$,

$$j + kr + s \in L \quad \text{iff} \quad j + r + s \in L.$$

In particular, if $m \geq p := j + r$, say $m = j + r + s$, then (take $k = 2$)

$$m + r \in L \quad \text{iff} \quad m \in L.$$

Now let L be ultimately periodic. Choose corresponding $p, r \in \mathbb{N}_+$. Set $L_1 := \{m \in L \mid m < p\}$ and $L_2 := \{m \in L \mid p \leq m < p + r\}$. Then, by periodicity, $L = L_1 \cup L_2 \cup \{m + kr \mid m \in L_2, k \geq 1\}$. So L is the union of the finite (and hence regular) sets L_1 and L_2 and of the languages denoted by the regular expressions $\mathbf{a}^m(\mathbf{a}^r)^+$ with $m \in L_2$. Thus L is regular. \square

As a consequence, the classes of finite ordered structures of vocabulary $\{<\}$ that are axiomatizable in MSO coincide with the ultimately periodic ones.

Example 6.3.3 For $\mathbb{A} = \{a, b\}$ the set

$$L := \{u \in \mathbb{A}^+ \mid \text{the number of } a\text{'s in } u \text{ equals the number of } b\text{'s in } u\}$$

is not accepted by an automaton. Otherwise, choose n and a representation vw^2x of $a^n b^n$ according to the Pumping Lemma. Since $|vw| \leq n$, we have $w \in \{a\}^+$. Hence the string vw^2x contains more a 's than b 's and therefore, $vw^2x \notin L$ (while $vw^2x = a^n b^n \in L$), a contradiction. \square

We use this example to prove nonaxiomatizability results for monadic second-order logic: A graph (G, E^G) is *bipartite*, if there is an $X \subseteq G$ such that $E^G \subseteq (X \times (G \setminus X)) \cup ((G \setminus X) \times X)$, and it is *balanced*, if the set X can be chosen such that, moreover, $\|X\| = \|G \setminus X\|$. Denote by BAL the class of finite balanced graphs and by $\text{BAL}_{<}$ the class of finite balanced graphs carrying an arbitrary ordering on their universe,

$$\text{BAL}_{<} := \{(G, <) \mid G \in \text{BAL}, < \text{ an ordering of } G\}.$$

Proposition 6.3.4 *The class $\text{BAL}_{<}$ – and hence the class BAL – is not axiomatizable in monadic second-order logic.*

Proof. Suppose that $\text{BAL}_{<} = \text{Mod}(\varphi)$ for a sentence φ of MSO. Let $\mathbb{A} = \{a, b\}$ and let L be as in the preceding example. For $u \in \mathbb{A}^+$ recall that $\mathcal{B}_u = (B_u, <, P_a, P_b)$ denotes a word model associated with u , say, with $B_u = \{1, \dots, |u|\}$ and $<$ the natural ordering. Let $\mathcal{G}_u = (B_u, R_u)$ be the bipartite graph given by

$$R_u := \{(i, j) \in B_u \times B_u \mid P_a i \text{ iff } P_b j\}.$$

Then,

$$(\mathcal{G}_u, <) \in \text{BAL}_{<} \quad \text{iff} \quad u \in L$$

and

$$(\mathcal{G}_u, <) \models \varphi \quad \text{iff} \quad \mathcal{B}_u \models \varphi \frac{(P_a \dots \leftrightarrow P_b _)}{E \dots _}$$

(where $\varphi \frac{(P_a \dots \leftrightarrow P_b _)}{E \dots _}$ is obtained by replacing any subformula in φ of the form Exy by $(P_a x \leftrightarrow P_b y)$). Therefore, $\text{Mod}(\varphi \frac{(P_a \dots \leftrightarrow P_b _)}{E \dots _}) = \{\mathcal{B}_u \mid u \in L\}$. Theorem 6.2.3 now implies that L is accepted by an automaton, which contradicts the preceding example. \square

Corollary 6.3.5 (a) *Let HAM be the class of finite graphs with a Hamiltonian circuit. Then, HAM and $\text{HAM}_{<}$ are not axiomatizable in MSO.*
 (b) *Let CHS be the set of finite graphs which contain a clique of at least half their size. Then, CHS and $\text{CHS}_{<}$ are not axiomatizable in MSO.*

Proof. (a) A graph of the form $(X \dot{\cup} Y, E)$ with $E = \{(a, b) \mid (a \in X, b \in Y) \text{ or } (a \in Y, b \in X)\}$ has a Hamiltonian circuit iff it is balanced. If $\text{HAM}_{<} = \text{Mod}(\varphi)$ for an MSO-sentence φ , then the sentence

$$\exists X (\forall x \forall y (Exy \rightarrow (Xx \leftrightarrow \neg Xy))) \wedge \varphi \frac{(X \dots \leftrightarrow \neg X _)}{E \dots _}$$

would axiomatize the class $\text{BAL}_{<}$.

(b) Suppose that $\text{CHS}_{<} = \text{Mod}(\varphi)$ for some φ of MSO. Then an axiomatization of $\text{BAL}_{<}$ in MSO would be given by

$$\begin{aligned} & \exists X (\forall x \forall y (Exy \rightarrow (Xx \leftrightarrow \neg Xy))) \\ & \wedge \varphi \frac{(X \dots \wedge X _ \wedge \neg \dots = _)}{E \dots _} \quad \wedge \quad \varphi \frac{(\neg X \dots \wedge \neg X _ \wedge \neg \dots = _)}{E \dots _} \end{aligned}$$

(note that the conjunction in the last line implies that both X and its complement have size at least half of the universe). \square

Exercise 6.3.6 Show that a graph \mathcal{G} is bipartite iff it contains no cycle of odd length (i.e., for no $n \geq 0$ there are a_1, \dots, a_{2n+1} such that $a_1 E^{\mathcal{G}} a_2, a_2 E^{\mathcal{G}} a_3, \dots, a_{2n} E^{\mathcal{G}} a_{2n+1}$, and $a_{2n+1} E^{\mathcal{G}} a_1$). Conclude that the class of bipartite graphs is not first-order axiomatizable (of course, it is axiomatizable in MSO). \square

6.4 First-Order Definability

We turn to the problem of characterizing the languages that are accepted by automata and are first-order definable. The proof of 6.2.1 shows that second-order quantifiers are only needed for the positive closure, that is, in the transition from a regular expression r to r^+ . Therefore, if r does not contain the symbol $^+$, the language L denoted by r is first-order definable. But a simple induction on the length of such an r shows that then L must be finite.

Example 6.4.1 Let \mathbb{A} be an alphabet. For $a \in \mathbb{A}$ the language $\mathbb{A}^+ \setminus \{a\}$ is infinite but first-order definable by

$$\varphi_W \wedge (\exists x \neg \psi_{\min}(x) \vee \exists x (\psi_{\min}(x) \wedge \neg P_a x))$$

(for the definition of φ_W and ψ_{\min} compare section 6.2). □

We thus see that the class of languages denoted by regular expressions without $^+$ is not closed under complementation whereas the class of first-order definable languages is. Therefore, we add closure under complementation in the definition of *plus free* regular expressions:³

- \emptyset , \mathbf{a} (for $a \in A$) are plus free regular expressions.
- If r and s are plus free regular expressions then so are $\sim r$, $(r \cup s)$, (rs) .

If r denotes the language L , then $\sim r$ denotes $\mathbb{A}^+ \setminus L$. A language is said to be *plus free regular* if it is denoted by a plus free regular expression.

Theorem 6.4.2 *A language is plus free regular iff it is definable in first-order logic.*

Proof. If the language L is definable by the first-order sentence φ , then $\mathbb{A}^+ \setminus L$ is definable by $(\varphi_W \wedge \neg \varphi)$. This observation together with the corresponding parts of the proof of 6.2.1 shows that any plus free regular language is first-order definable.

We turn to the other direction. Recall that $\tau(\mathbb{A}) = \{<\} \cup \{P_a \mid a \in \mathbb{A}\}$. For convenience, we add a constant \min to this vocabulary, which henceforth will always denote the first element; more precisely, we shall only look at models of $\varphi_W \wedge \psi_{\min}(\min)$. By induction on the quantifier rank of the FO $[\tau(\mathbb{A}) \cup \{\min\}]$ -sentence φ we show for a language L that if

$$\text{Mod}(\varphi_W \wedge \psi_{\min}(\min) \wedge \varphi) = \{(B_u, \min^{B_u}) \mid u \in L\}$$

then L is plus free regular. (This gives the claim.)

First assume that φ is atomic. Then φ is $\min = \min$ or $P_a \min$ for some $a \in \mathbb{A}$. In the first case, L is \mathbb{A}^+ , which can be denoted by the plus free

³ In the literature languages and, in particular, languages accepted by automata, may contain the empty word. Then the operation $^+$ is replaced by the operation * , where $L^* := L^+ \cup \{\lambda\}$ and the role of ‘plus free’ is taken over by ‘star free’.

regular expression $\sim \emptyset$. Let φ be $P_a \text{ min}$. Then $L = \{a\} \cup \{a\} \mathbb{A}^+$. Therefore, L is denoted by $\mathbf{a} \cup \mathbf{a}(\sim \emptyset)$.

If the languages defined by the sentences φ and ψ are denoted by the plus free expressions r and s , then $\sim r$ and $r \cup s$ correspond to the sentences $\neg\varphi$ and $(\varphi \vee \psi)$, respectively.

Let $\varphi = \exists x\psi(x)$. Then $\text{Mod}(\varphi_W \wedge \psi_{\min}(\min) \wedge \exists x\psi(x)) =$

$$(*) \quad \begin{aligned} & \text{Mod}(\varphi_W \wedge \psi_{\min}(\min) \wedge \psi(\min)) \\ \cup & \text{Mod}(\varphi_W \wedge \psi_{\min}(\min) \wedge \exists x(\neg x = \min \wedge \psi(x))). \end{aligned}$$

By induction hypothesis, the first class of structures in $(*)$ corresponds to a plus free regular language. Concerning the second class, let c be a new constant. Then the finite models of $\varphi_W \wedge \psi_{\min}(\min) \wedge \exists x(\neg x = \min \wedge \psi(x))$ are the $\tau(\mathbb{A}) \cup \{\min\}$ -reducts of the finite structures $(\mathcal{A}, \min^A, c^A)$ such that

$$(\mathcal{A}, \min^A, c^A) \models \varphi_W \wedge \psi_{\min}(\min) \wedge \neg c = \min \wedge \psi(c).$$

Note that any such structure can be written in the form

$$(\mathcal{A}, \min^A, c^A) = (\mathcal{A}_1 \triangleleft \mathcal{A}_2, \min^A, c^A),$$

where \triangleleft denotes the ordered sum and where $(\mathcal{A}_1, \min^A) \models (\varphi_W \wedge \psi_{\min}(\min))$ and $(\mathcal{A}_2, c^A) \models (\varphi_W \wedge \psi_{\min}(c))$. Let m be the quantifier rank of ψ . Choose the – up to logical equivalence – finite set $\{(\psi_i(\min), \chi_i(c)) \mid i \in I\}$ of pairs of FO-sentences of quantifier rank $\leq m$ such that

$$\begin{aligned} & (\mathcal{A}_1, \min^{A_1}) \models (\varphi_W \wedge \psi_{\min}(\min) \wedge \psi_i(\min)) \\ \text{and} & \quad (\mathcal{A}_2, c^{A_2}) \models (\varphi_W \wedge \psi_{\min}(c) \wedge \chi_i(c)) \\ \text{imply} & \quad (\mathcal{A}_1, \min^{A_1}) \triangleleft (\mathcal{A}_2, c^{A_2}) \models \psi(c). \end{aligned}$$

By induction hypothesis there are plus free regular expressions r_i and s_i denoting the languages defined by $\varphi_W \wedge \psi_{\min}(\min) \wedge \psi_i(\min)$ and $\varphi_W \wedge \psi_{\min}(\min) \wedge \chi_i(\min)$, respectively. Then the plus free regular expression $\bigcup_{i \in I} (r_i s_i)$ denotes the language defined by $(\varphi_W \wedge \psi_{\min}(\min) \wedge \exists x(\neg x = \min \wedge \psi(x)))$. Note that, if $(\mathcal{A}_1 \triangleleft \mathcal{A}_2, \min^{A_1}, c^{A_2}) \models \psi(c)$ then (by 2.3.11) the pair $(\varphi_{(\mathcal{A}_1, \min^{A_1})}^m, \varphi_{(\mathcal{A}_2, c^{A_2})}^m)$ of m -isomorphism types belongs (up to logical equivalence) to $\{(\psi_i(\min), \chi_i(c)) \mid i \in I\}$. \square

Again, for $\mathbb{A} = \{a\}$ identify \mathbb{A}^+ with the set \mathbb{N}_+ of positive natural numbers. Examples 6.3.2 and 6.3.3 show that automata do not have the ability to count; for instance, they cannot recognize if a given string has prime length; more explicitly, the set $\{p \mid p \text{ a prime}\}$ is not accepted by an automaton. On the other hand, we saw in Example 6.3.2 that automata are capable to count modulo a natural number, e.g. the set $\{5n \mid n \geq 1\}$ is accepted by an automaton. But first-order logic even lacks this restricted counting ability;

in fact, it is an immediate consequence of 2.3.6 that a subset L of \mathbb{N}_+ is first-order definable iff for some $n \geq 1$

$$\{m \mid m \geq n\} \cap L = \emptyset \quad \text{or} \quad \{m \mid m \geq n\} \subseteq L.$$

Before we state the corresponding theorem for word models over *arbitrary* alphabets, we remark that we encounter this noncounting feature of first-order logic at various points in the book.

Theorem 6.4.3 *For a language $L \subseteq \mathbb{A}^+$ accepted by an automaton the following are equivalent:*

- (i) L is definable in first-order logic.
- (ii) L is noncounting in the sense that there is an integer $k \geq 1$ such that for every $y \in \mathbb{A}^+$ and $x, z \in \mathbb{A}^*$,

$$xy^kz \in L \quad \text{iff} \quad xy^{k+1}z \in L.$$

Proof. We only prove the implication from (i) to (ii) and refer the reader to the literature for the involved proof of the other half (cf. 6.4.4). Suppose $\{\mathcal{B}_u \mid u \in L\} = \text{Mod}(\varphi)$ for $\varphi \in \text{FO}[\tau(\mathbb{A})]$. Let $k := 2^m + 1$, where m is the quantifier rank of φ . Then by 2.3.13, for any $y \in \mathbb{A}^+$ we have

$$\mathcal{B}_{y^k} \cong \triangleleft^k \mathcal{B}_y \equiv_m \triangleleft^{k+1} \mathcal{B}_y \cong \mathcal{B}_{y^{k+1}}.$$

Using 2.3.10, we obtain

$$\mathcal{B}_{xy^kz} \cong \mathcal{B}_x \triangleleft \mathcal{B}_{y^k} \triangleleft \mathcal{B}_z \equiv_m \mathcal{B}_x \triangleleft \mathcal{B}_{y^{k+1}} \triangleleft \mathcal{B}_z \cong \mathcal{B}_{xy^{k+1}z}.$$

In particular,

$$\mathcal{B}_{xy^kz} \models \varphi \quad \text{iff} \quad \mathcal{B}_{xy^{k+1}z} \models \varphi$$

and hence,

$$xy^kz \in L \quad \text{iff} \quad xy^{k+1}z \in L. \quad \square$$

As the results of this section show, the plus operation cannot be captured in first-order logic. An instance of this operation can be viewed as the fixed-point of a monotone operation. In fact, let $L \subseteq \mathbb{A}^+$ be a language. Define $C_L : \text{Pow}(\mathbb{A}^*) \rightarrow \text{Pow}(\mathbb{A}^*)$ by

$$C_L(M) := L \cup ML.$$

Then

- (a) C_L is monotone, that is, $M_1 \subseteq M_2$ implies $C_L(M_1) \subseteq C_L(M_2)$.
- (b) For $n \geq 1$, $\underbrace{C_L(\dots(C_L(\emptyset))\dots)}_{n\text{-times}} = L \cup L^2 \cup \dots \cup L^n$.

M is a *fixed-point* of C_L if $C_L(M) = M$. It can easily be proved that the least – with respect to set-theoretical inclusion – fixed-point of C_L is given by

$$C_L(\emptyset) \cup C_L(C_L(\emptyset)) \cup C_L(C_L(C_L(\emptyset))) \cup \dots$$

Hence by (b),

$$L^+ \text{ is the least fixed-point of } C_L.$$

Fixed-points of monotone operations play a prominent role in the general theory. We shall study them in Chapters 7 and 8. In particular, we shall see that polynomial time queries on structures with an underlying ordering can be captured in a logic extending first-order logic by adding the ability to express fixed-points of definable monotone operations.

Notes 6.4.4 The connection between finite automata and monadic second-order logic as expressed in Theorem 6.2.3 is due to [15, 140]. For a proof of Theorem 6.4.3 and references on finite automata compare [136] or [138]. Some examples are taken from [17].

7. Descriptive Complexity Theory

In Chapter 1 we gave the example of a database D that contains the names of the main cities in the world and the pairs (a, b) of cities such that a given airline offers service from a to b without stopover. D may be interpreted as a first-order structure, more precisely, as a digraph $\mathcal{G} = (G, E^G)$, where G is the set of cities and $E^G ab$ means that there is a flight from a to b without stopover. Now, first-order logic may be viewed as a query language. For example, let

$$\varphi(x, y) := Exy \vee \exists z(Exz \wedge Ezy).$$

If φ is thought of as a query to D , the response will consist in the set of pairs (a, b) of cities such that a can be reached from b with at most one stop. First-order logic provides a rich class of database queries. However, there are plausible queries which are not first-order expressible. The query Q “Is it possible to fly from x to y using only the airline in question?” to databases of type D is an example. Also from a computational point of view answers to Q are more complex than answers to φ . Descriptive complexity theory analyzes the complexity of all queries definable in a given logic, the central question being the following: Given a complexity class \mathcal{C} , is there a logic \mathcal{L} such that the queries definable in \mathcal{L} are precisely the queries in \mathcal{C} ? For short: Given a complexity class \mathcal{C} , is there a logic that captures \mathcal{C} ? In the positive case, there mostly will be an effective procedure translating every formula of such a logic into a program for the corresponding query. In this sense the logic can be viewed as a higher programming language for \mathcal{C} .

In the first section we briefly introduce some extensions of first-order logic which have been considered with this aim in mind. Their model theory will be studied in Chapter 8. The second section is devoted to complexity classes of structures. In particular, we describe how structures may be viewed as inputs to Turing machines. In section 7.3 we show how certain logics can be used to describe computations of a given complexity class. In section 7.4 we look at the other direction and study the complexity of the satisfaction relation of such logics. By bringing together both sides we are lead to the descriptive characterizations of complexity classes we aim at, thus constituting a bridge between complexity theory and model theory. First consequences are given in the last section of this chapter.

Throughout this chapter all structures will be finite.

7.1 Some Extensions of First-Order Logic

We introduce inflationary fixed-point logic, partial fixed-point logic, deterministic transitive closure logic, and transitive closure logic. These logics are obtained from first-order logic by adding operations well-suited to describe iterative and recursive procedures, for example, the behaviour of computing machines. As already remarked in the introduction to this chapter, their model theory will be studied in Chapter 8.

Let M be a finite nonempty set. Denote by $\text{Pow}(M)$ the power set of M . A function $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ induces a sequence

$$\emptyset, F(\emptyset), F(F(\emptyset)), \dots$$

of subsets of M . For its members we write F_0, F_1, \dots . So $F_0 = \emptyset$ and $F_{n+1} = F(F_n)$. Suppose there is an $n_0 \geq 0$ such that $F_{n_0+1} = F_{n_0}$, that is, $F(F_{n_0}) = F_{n_0}$. Then $F_m = F_{n_0}$ for all $m \geq n_0$. We denote F_{n_0} by F_∞ and say that the *fixed-point* F_∞ of F *exists*. In case the fixed-point F_∞ does not exist, we agree to set $F_\infty := \emptyset$.

F is said to be *inflationary* if for all $X \subseteq M$,

$$X \subseteq F(X).$$

Lemma 7.1.1 (a) *The sequence $(F_n)_{n \geq 0}$ is periodic, more precisely: There are $m < 2^{\|M\|}$ and $l \geq 1$ such that $F_k = F_{k+l}$ for all $k \geq m$.*

(b) *If F_∞ exists then $F_\infty = F_{2^{\|M\|} - 1}$.*

(c) *If F is inflationary then F_∞ exists and $F_\infty = F_{\|M\|}$.*

Proof. (a) As $\text{Pow}(M)$ has $2^{\|M\|}$ elements, there are $m < 2^{\|M\|}$ and $l \geq 1$ such that $F_m = F_{m+l}$. Therefore, $F_{m+1} = F(F_m) = F(F_{m+l}) = F_{m+1+l}$, $F_{m+2} = F_{m+2+l}, \dots$

(b) Choose $m < 2^{\|M\|}$ and $l \geq 1$ according to (a). If $F_m = F_{m+1}$ then $F_m = F_{2^{\|M\|} - 1} = F_\infty$. If $F_m \neq F_{m+1}$ then, by (a), for $s \geq m$, we get $(F_m =) F_{m+s \cdot l} \neq F_{m+1+s \cdot l} (= F_{m+1})$ and hence, F_∞ does not exist.

(c) By assumption, $F_0 \subseteq F_1 \subseteq \dots \subseteq M$. Since M has $\|M\|$ elements, this sequence must get constant not later than with $F_{\|M\|}$. \square

Let $\varphi(x_1, \dots, x_k, \bar{u}, X, \bar{Y})$ be a formula in the vocabulary τ , where the relation variable X has arity k ; moreover, let \mathcal{A} be a τ -structure, \bar{b} an interpretation of \bar{u} in \mathcal{A} , and \bar{S} an interpretation of \bar{Y} over \mathcal{A} . Then $\varphi, \mathcal{A}, \bar{b}$, and \bar{S} give rise to an operation $F^\varphi : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$ defined by

$$F^\varphi(R) := \{(a_1, \dots, a_k) \mid \mathcal{A} \models \varphi[a_1, \dots, a_k, \bar{b}, R, \bar{S}]\}.$$

(The notation F^φ does not make explicit all relevant data.)

To give an example, let $\mathcal{G} = (G, E^G)$ be a graph and

$$\varphi_0(x, y, X) := (Exy \vee \exists z(Xxz \wedge Ezy))$$

with xy corresponding to \bar{x} above. Then, $F_0^{\varphi_0} = \emptyset$, $F_1^{\varphi_0} = F^{\varphi_0}(\emptyset) = E^G$, $F_2^{\varphi_0} = F^{\varphi_0}(E^G) = E^G \cup \{(a, b) \mid (E^G ac \text{ and } E^G cb) \text{ for some } c \in G\}$. By induction on n one shows that

$$F_n^{\varphi_0} = \{(a, b) \mid \text{there is a path of length } \leq n \text{ from } a \text{ to } b\},$$

and hence,

$$F_\infty^{\varphi_0} = \{(a, b) \mid \text{there is a path from } a \text{ to } b\}.$$

Note that for φ as above the function $F^{(X\bar{x}\vee\varphi)}$ (where $\bar{x} = x_1 \dots x_k$) is inflationary.

One obtains *Inflationary Fixed-Point Logic* FO(IFP) and *Partial Fixed-Point Logic* FO(PFP) by closing first-order logic FO under inflationary and arbitrary fixed-points of definable operations, respectively. We state the precise definitions.

For a vocabulary τ the class FO(IFP)[τ] of formulas of FO(IFP) of vocabulary τ is given by the calculus (we use the succinct notation

$$\frac{\varphi_1, \dots, \varphi_n}{\varphi}$$

for the clause “If $\varphi_1, \dots, \varphi_n$ are formulas then φ is a formula”)

- $\frac{}{\varphi}$ where φ is an atomic second-order formula over τ
- $\frac{\varphi}{\neg\varphi}$; $\frac{\varphi, \psi}{(\varphi \vee \psi)}$; $\frac{\varphi}{\exists x\varphi}$
- $\frac{\varphi}{[\text{IFP}_{\bar{x}, X}\varphi]\bar{t}}$ where the lengths of \bar{x} and \bar{t} are the same and coincide with the arity of X .

For FO(PFP) we replace the last rule by

- $\frac{\varphi}{[\text{PFP}_{\bar{x}, X}\varphi]\bar{t}}$ where the lengths of \bar{x} and \bar{t} are the same and coincide with the arity of X .

Sentences are formulas without free first-order and second-order variables, where the free occurrence of variables is defined in the standard way, adding, for example, for FO(IFP) the clause

$$\text{free}([\text{IFP}_{\bar{x}, X}\varphi]\bar{t}) := \text{free}(\bar{t}) \cup (\text{free}(\varphi) \setminus \{\bar{x}, X\}).$$

The semantics is defined inductively w.r.t. the calculus above,

$$[\text{IFP}_{\bar{x}, X}\varphi]\bar{t} \text{ meaning that } \bar{t} \in F_\infty^{(X\bar{x}\vee\varphi)}$$

and

$$[\text{PFP}_{\bar{x}, X} \varphi] \bar{t} \text{ meaning that } \bar{t} \in F_{\infty}^{\varphi}.$$

More precisely: If X is k -ary and if the variables free in $[\text{IFP}_{\bar{x}, X} \varphi] \bar{t}$ are among \bar{u} and \bar{Y} , and \bar{b} and \bar{S} are interpretations in \mathcal{A} of \bar{u} and \bar{Y} , respectively, then

$$\begin{aligned} \mathcal{A} \models [\text{IFP}_{\bar{x}, X} \varphi] \bar{t} [\bar{b}, \bar{S}] & \quad \text{iff} \quad (t_1[\bar{b}], \dots, t_k[\bar{b}]) \in F_{\infty}^{(X\bar{x} \vee \varphi)}, \\ \mathcal{A} \models [\text{PFP}_{\bar{x}, X} \varphi] \bar{t} [\bar{b}, \bar{S}] & \quad \text{iff} \quad (t_1[\bar{b}], \dots, t_k[\bar{b}]) \in F_{\infty}^{\varphi}. \end{aligned}$$

Examples 7.1.2 (a) In the language of graphs the formula

$$\psi_0(x, y) := [\text{IFP}_{xy, X} (Exy \vee \exists z (Xxz \wedge Ezy))] xy$$

of FO(IFP) expresses that x, y are connected by a path. Hence, the class CONN of connected graphs is axiomatizable in FO(IFP) by $\forall x \forall y (\neg x = y \rightarrow \psi_0(x, y))$ (and the graph axioms). It is not axiomatizable in FO (cf. 2.3.8).

(b) For $\tau = \{<, S, \text{min}, \text{max}\}$ the sentence

$$\neg[\text{IFP}_{x, X} (x = \text{min} \vee \exists y \exists z (Xy \wedge Syz \wedge Szx))] \text{max}$$

of FO(IFP) together with the ordering axioms axiomatizes the class of orderings of even cardinality. The same holds for the class of orderings of odd cardinality and the FO(PFP)-sentence $\exists x [\text{PFP}_{x, X} \psi(x, X)] x$ where

$$\psi(x, X) = (\forall y \neg Xy \wedge x = \text{min}) \vee (X \text{max} \wedge x = \text{max}) \vee \exists y (Xy \wedge \exists u (Syu \wedge Suu)). \quad \square$$

To compare the expressive power of logics we introduce the following relations.

Definition 7.1.3 Let \mathcal{L}_1 and \mathcal{L}_2 be logics.

- (a) $\mathcal{L}_1 \leq \mathcal{L}_2$ (read: \mathcal{L}_1 is at most as expressive as \mathcal{L}_2) if for every τ and every sentence $\varphi \in \mathcal{L}_1[\tau]$ there is a sentence $\psi \in \mathcal{L}_2[\tau]$ such that $\text{Mod}(\varphi) = \text{Mod}(\psi)$.¹
- (b) $\mathcal{L}_1 \equiv \mathcal{L}_2$ (read: \mathcal{L}_1 and \mathcal{L}_2 have the same expressive power) if $\mathcal{L}_1 \leq \mathcal{L}_2$ and $\mathcal{L}_2 \leq \mathcal{L}_1$.
- (c) $\mathcal{L}_1 < \mathcal{L}_2$ if $\mathcal{L}_1 \leq \mathcal{L}_2$ and not $\mathcal{L}_2 \leq \mathcal{L}_1$. □

In most cases, a proof of $\mathcal{L}_1 \leq \mathcal{L}_2$ even yields that every formula of \mathcal{L}_1 is equivalent to a formula of \mathcal{L}_2 . In particular, $\mathcal{L}_1 \leq \mathcal{L}_2$ implies that $\mathcal{L}_1 \leq \mathcal{L}_2$ holds for all formulas of \mathcal{L}_1 containing only free *individual* variables (one simply replaces these variables by new constants). Thus, $\mathcal{L}_1 \leq \mathcal{L}_2$ will imply that every global relation definable in \mathcal{L}_1 is also definable in \mathcal{L}_2 .

By Example 7.1.2(a), we have $\text{FO} < \text{FO}(\text{IFP})$.

¹ $\mathcal{L}[\tau]$ denotes the class of formulas of \mathcal{L} of vocabulary τ .

Proposition 7.1.4 $\text{FO}(\text{IFP}) \leq \text{FO}(\text{PFP})$.

Proof. Note that $[\text{IFP}_{\bar{x}, X} \varphi] \bar{t}$ is equivalent to $[\text{PFP}_{\bar{x}, X} (X\bar{x} \vee \varphi)] \bar{t}$. \square

Let R be a binary relation on a set M , $R \subseteq M^2$. The *transitive closure* $\text{TC}(R)$ of R is defined by

$$\text{TC}(R) := \{(a, b) \in M^2 \mid \text{there exist } n > 0 \text{ and } e_0, \dots, e_n \in M \text{ such that } a = e_0, b = e_n, \text{ and for all } i < n, (e_i, e_{i+1}) \in R\}.$$

And the *deterministic transitive closure* $\text{DTC}(R)$ is defined by

$$\text{DTC}(R) := \{(a, b) \in M^2 \mid \text{there exist } n > 0 \text{ and } e_0, \dots, e_n \in M \text{ such that } a = e_0, b = e_n, \text{ and for all } i < n, e_{i+1} \text{ is the unique } e \text{ for which } (e_i, e) \in R\}.$$

Transitive Closure Logic $\text{FO}(\text{TC})$ and *Deterministic Transitive Closure Logic* $\text{FO}(\text{DTC})$ are obtained by closing FO under the transitive closure and the deterministic transitive closure of definable relations, respectively. More precisely: For a vocabulary τ the class of formulas of $\text{FO}(\text{TC})$ of vocabulary τ is given by the calculus

- $\frac{}{\varphi}$ where φ is an atomic first-order formula over τ
- $\frac{\varphi}{\neg\varphi}$; $\frac{\varphi, \psi}{(\varphi \vee \psi)}$; $\frac{\varphi}{\exists x\varphi}$
- $\frac{\varphi}{[\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s}\bar{t}}$ where the variables in $\bar{x}\bar{y}$ are pairwise distinct and where the tuples $\bar{x}, \bar{y}, \bar{s}$, and \bar{t} are all of the same length, \bar{s} and \bar{t} being tuples of terms.

For $\text{FO}(\text{DTC})$ the last rule is replaced by $\frac{\varphi}{[\text{DTC}_{\bar{x}, \bar{y}} \varphi] \bar{s}\bar{t}}$ with the same side conditions. We define

$$\text{free}([\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s}\bar{t}) := \text{free}(\bar{s}) \cup \text{free}(\bar{t}) \cup (\text{free}(\varphi) \setminus \{\bar{x}, \bar{y}\}),$$

and similarly for $\text{FO}(\text{DTC})$.

The meaning of $[\text{TC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{u})] \bar{s}\bar{t}$ is

$$(\bar{s}, \bar{t}) \in \text{TC}(\{(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{y}, \bar{u})\}),$$

and the meaning of $[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{u})] \bar{s}\bar{t}$ is

$$(\bar{s}, \bar{t}) \in \text{DTC}(\{(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{y}, \bar{u})\})$$

(here $\{(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{y}, \bar{u})\}$ is considered as a binary relation on the set of $\text{length}(\bar{x})$ -tuples of the universe).

Thus a graph is connected if it is a model of $\forall x \forall y (\neg x = y \rightarrow [\text{TC}_{x, y} E x y] x y)$.

Proposition 7.1.5 (a) $\text{FO}(\text{DTC}) \leq \text{FO}(\text{TC})$.

(b) $\text{FO}(\text{TC}) \leq \text{FO}(\text{IFP})$.

Proof. For (a) note that

$$\models_{\text{fin}} [\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{u})] \bar{s} \bar{t} \leftrightarrow [\text{TC}_{\bar{x}, \bar{y}}(\varphi(\bar{x}, \bar{y}, \bar{u}) \wedge \forall \bar{z}(\varphi(\bar{x}, \bar{z}, \bar{u}) \rightarrow \bar{z} = \bar{y}))] \bar{s} \bar{t}$$

and for (b) that

$$\models_{\text{fin}} [\text{TC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{u})] \bar{s} \bar{t} \leftrightarrow [\text{IFP}_{\bar{x}, \bar{y}, X}(\varphi(\bar{x}, \bar{y}, \bar{u}) \vee \exists \bar{v}(X \bar{x} \bar{v} \wedge \varphi(\bar{v}, \bar{y}, \bar{u})))] \bar{s} \bar{t}.$$

□

In the following, axiomatizability in a logic (in the sense of the following definition) will be a major issue.

Definition 7.1.6 Let K be a class of τ -structures and \mathcal{L} a logic. K is *axiomatizable* in \mathcal{L} , if there is a sentence of \mathcal{L} of vocabulary τ such that $K = \text{Mod}(\varphi)$.

□

Sometimes, when relating logics and complexity classes it is convenient to restrict oneself to sufficiently large structures. This does not affect problems of axiomatizability. In fact, for a class K of structures and $m \geq 1$ denote by K_m the subclass of K of structures of cardinality $\geq m$,

$$K_m := \{\mathcal{A} \mid \mathcal{A} \in K, \|\mathcal{A}\| \geq m\}.$$

For every finite structure \mathcal{A} there is a sentence $\varphi_{\mathcal{A}}$ of FO characterizing \mathcal{A} up to isomorphism (cf. 2.1.1), that is,

$$\mathcal{B} \models \varphi_{\mathcal{A}} \text{ iff } \mathcal{B} \cong \mathcal{A}$$

holds for all \mathcal{B} . Hence, for any logic \mathcal{L} with $\text{FO} \leq \mathcal{L}$,

$$K \text{ is axiomatizable in } \mathcal{L} \text{ iff } K_m \text{ is axiomatizable in } \mathcal{L}.$$

In fact, setting $\varphi_m := \bigvee \{\varphi_{\mathcal{A}} \mid \mathcal{A} \in K, \|\mathcal{A}\| < m\}$ we have that $K = \text{Mod}(\varphi)$ implies $K_m = \text{Mod}(\varphi \wedge \neg \varphi_m)$, and $K_m = \text{Mod}(\psi)$ implies $K = \text{Mod}(\psi \vee \varphi_m)$.

7.2 Turing Machines and Complexity Classes

The aim of the present section is to recall basic definitions and results from computation theory, to fix our computation model for structures as inputs, and to introduce the corresponding complexity classes. It is not intended as a substitute for a course in computation theory, and some familiarity with the basic notions will be helpful. For references compare the notes 7.5.26.

We adopt a computation model which belongs to the most popular ones in theoretical computer science, the Turing machine model. Our choice is

motivated by the fact that Turing machine computations allow for simple descriptions and for natural definitions of complexity classes. However, as Turing machines deal only with strings, any other data must be coded by strings. The tedious work of encoding is one price we have to pay for the advantages. Moreover, the model does not well reflect parallel, distributed, or real time computations, as only local changes are made in a step. However, these shortages are not really important for our aims.

For the following, we fix a finite alphabet \mathbb{A} . A *Turing machine* M is a finite device that performs operations on a tape which is bounded to the left and unbounded to the right and divided into *squares* (or *cells*). The machine operates stepwise, each step leading from one situation to a new one. In any situation every square of the tape either contains a single symbol from \mathbb{A} or is blank. In the latter case we say that it contains the symbol “blank”. There is one exception: the leftmost or “virtual” cell always contains an endmark, the “virtual” letter α (which is not in \mathbb{A}). M has a head which, in any situation, scans a single square of the tape and, in any step of a computation, erases or replaces the scanned symbol by another one and moves one cell to the left or to the right or remains at its place (we speak of a *read-and-write head*).

In every situation, M is in one of the states of a finite set $\text{State}(M)$, the *set of states* of M . $\text{State}(M)$ contains a special state s_0 , the *initial* state, and special states s_+ , the *accepting* state, and s_- , the *rejecting* state. We assume that s_0, s_+ , and s_- are pairwise distinct. The action or behaviour of M in a situation depends on the current state of M and on the symbol currently being scanned by the head. It is given by $\text{Instr}(M)$, the *set of instructions* of M . Each instruction has the form

$$(*) \quad sa \rightarrow s'bh$$


where

- $s, s' \in \text{State}(M)$, $s \neq s_+$, $s \neq s_-$
- $a, b \in \mathbb{A} \cup \{\alpha, \text{blank}\}$ and ($a = \alpha$ iff $b = \alpha$)
- $h \in \{-1, 0, 1\}$, and if $a = \alpha$ then $h \neq -1$.

The instruction $(*)$ means: If you are in state s and your head scans a cell with symbol a , replace a by b , move your head one cell to the left ($h = -1$), or to the right ($h = 1$), or don't move ($h = 0$); finally, change to state s' .

A machine M is *deterministic*, if for all $s \in S$ and $a \in \mathbb{A} \cup \{\alpha, \text{blank}\}$ there is at most one instruction of the form $(*)$ in $\text{Instr}(M)$. In order to emphasize that a machine is not required to be deterministic we sometimes call it *nondeterministic*.

As usual, denote by \mathbb{A}^* the set of words over \mathbb{A} and by \mathbb{A}^+ the set of nonempty words over \mathbb{A} . Let $u \in \mathbb{A}^*$, $u = a_1 \dots a_r$ with $a_i \in \mathbb{A}$. M is *started* with u if M begins a computation (or *run*) in state s_0 in the situation



The computation proceeds stepwise, each step corresponding to the execution of one instruction of M . The machine stops when it is in a state s scanning a symbol $a \in \mathbb{A} \cup \{\alpha, \text{blank}\}$ such that there is no instruction of the form $(*)$ in $\text{Instr}(M)$. If then $s = s_+$ we speak of an *accepting run*, and if $s = s_-$ of a *rejecting run*. M *accepts* u if there is at least one accepting run of M started with u , and M *rejects* u if all runs started with u are finite and rejecting.

Subsets of \mathbb{A}^+ are called *languages*. A language $L \subseteq \mathbb{A}^+$ is *accepted* by M if for all $u \in \mathbb{A}^+$,

$$M \text{ accepts } u \quad \text{iff} \quad u \in L.$$

L is *decided* by M if, in addition,

$$M \text{ rejects } u \quad \text{iff} \quad u \notin L.$$

Clearly, if M decides L then M accepts L . L is said to be *decidable* if it is decided by some deterministic Turing machine, and *acceptable* or *enumerable* if it is accepted by some nondeterministic Turing machine.

For a function $f : \mathbb{N} \rightarrow \mathbb{N}$ we say that M is *f time-bounded*, if for all $u \in \mathbb{A}^+$ accepted by M there is an accepting run of M started with u which has length at most $f(|u|)$ (recall that $|u|$ denotes the length of the word u). And M is *f space-bounded*, if for all $u \in \mathbb{A}^+$ accepted by M there is an accepting run which uses at most $f(|u|)$ cells before stopping.

Denote by $\mathbb{N}[x]$ the set of polynomials with coefficients from \mathbb{N} . A language $L \subseteq \mathbb{A}^+$ is in PTIME (“polynomial time”) or in PSPACE (“polynomial space”), if it is accepted by a deterministic machine that is p time-bounded or p space-bounded, respectively, for some polynomial $p \in \mathbb{N}[x]$. The classes NPTIME (“nondeterministic polynomial time”) and NPSPACE (“nondeterministic polynomial space”) are defined similarly, now allowing nondeterministic machines.

Immediately from the definitions one gets

$$\text{PTIME} \subseteq \text{NPTIME} \quad \text{and} \quad \text{PTIME} \subseteq \text{PSPACE} \subseteq \text{NPSPACE},$$

and one can show that

$$\text{NPTIME} \subseteq \text{PSPACE} \quad \text{and} \quad \text{PSPACE} = \text{NPSPACE}.$$

Hence,

$$\text{PTIME} \subseteq \text{NPTIME} \subseteq \text{PSPACE} (= \text{NPSPACE}).$$

Let $\mathbb{A} = \{a, b\}$ and $L := \{u \in \mathbb{A}^+ \mid u \text{ contains an even number of } a\text{'s}\}$. One can easily design a machine that accepts L and is time-bounded by the polynomial $x + 2$: the head just runs over the string, the state being “even” or “odd” depending on whether the number of a ’s already scanned is even or odd, respectively. Essentially, this machine does not need any “working space” but only the space for the input. In order to measure only the working space and to introduce complexity classes like LOGSPACE and NLOGSPACE, where the working space needed is smaller than the input space, it is convenient to separate the input and to introduce machines with an input tape and a work tape. For other purposes it might be useful to introduce several input tapes and several work tapes, maybe also (if calculating a function, for example) one or more output tapes, with a head for each tape, where the heads can move independently of each other. It turns out that the definition of the usual complexity classes does not depend on the number of tapes or on other peculiarities such as the form of the tape, i.e., whether it is unbounded to both sides or not. We shall often use this robustness, choosing, for example, the number of input and work tapes according to needs and to convenience.

7.2.1 Digression: Trahtenbrot’s Theorem

We have already remarked that first-order logic, when restricted to the finite, does not admit a complete proof calculus. In the following we give a proof of Trahtenbrot’s Theorem which immediately implies this fact.

Fix an alphabet \mathbb{A} . It is well-known that it is not decidable whether a deterministic Turing machine M accepts the empty word, by short, whether M halts (“Undecidability of the Halting Problem”).² We use this result to show for a certain vocabulary $\sigma(\mathbb{A})$:

Theorem 7.2.1 (Trahtenbrot’s Theorem) *Finite Satisfiability is not decidable, that is, the set*

$$\text{Sat}[\sigma(\mathbb{A})] := \{\varphi \mid \varphi \text{ is a sentence of FO}[\sigma(\mathbb{A})] \text{ satisfiable in the finite}\}$$

is not decidable.

Proof. We assign, in an effective way, to every deterministic machine M over \mathbb{A} a sentence φ_M of $\text{FO}[\sigma(\mathbb{A})]$ such that

$$(*) \quad \varphi_M \text{ is satisfiable in the finite iff } M \text{ halts.}$$

This gives the claim by the result just mentioned. Without loss of generality we may restrict ourselves to deterministic Turing machines M whose set of states is an initial segment $\{0, \dots, s_M\}$ of the natural numbers, $s_0 := 0$ being the initial state and 1 the accepting state, and which stop only in the accepting state. We number the cells of the tape as indicated by

² More exactly: One encodes Turing machines M in some natural way as words w_M . Then the language $\{w_M \mid M \text{ halts}\}$ is not decidable.

We leave it to the reader to write down the sentence $\varphi_{sa \rightarrow s'bh}$ for $h = -1$ and $h = 1$.

- (4) $\exists t$ State 1 t
(the accepting state is reached). □

Coding $\sigma(\mathbb{A})$ -structures by graphs (cf. 11.2.5), one obtains from the undecidability of $\text{Sat}[\sigma(\mathbb{A})]$, that for a binary relation symbol E the set

$\text{Graph-Sat} := \{\varphi \mid \varphi \text{ is an FO}\{\{E\}\}\text{-sentence satisfiable in a finite graph}\}$

is not decidable. Hence, for any vocabulary τ containing an at least binary relation symbol, the set

$\text{Sat}[\tau] := \{\varphi \mid \varphi \text{ is an FO}[\tau]\text{-sentence satisfiable in the finite}\}$

is not decidable.

Clearly, there is a decision procedure that, for any finite structure \mathcal{A} whose universe is a set of natural numbers and for any first-order sentence φ , checks whether $\mathcal{A} \models \varphi$. Since φ is satisfiable in the finite iff there is a model \mathcal{A} of φ with $A = \{0, \dots, n\}$ for some n , we can use this decision procedure to enumerate $\text{Sat}[\tau]$. Together with Trahtenbrot's Theorem this shows:

Theorem 7.2.2 *If τ contains an at least binary relation symbol then the set*

$\text{Val}[\tau] := \{\varphi \in \text{FO}[\tau] \mid \varphi \text{ is a sentence valid in all finite structures}\}$

of sentences valid in all finite structures is not enumerable.

Proof. For any sentence φ we have

$$(+) \quad \varphi \notin \text{Sat}[\tau] \quad \text{iff} \quad \neg\varphi \in \text{Val}[\tau].$$

If $\text{Val}[\tau]$ would be enumerable, then by (+), $\text{FO}[\tau] \setminus \text{Sat}[\tau]$ would be enumerable, too. This would lead to a decision procedure for $\text{Sat}[\tau]$ (contradicting Trahtenbrot's Theorem): Given a sentence φ , start enumeration procedures for $\text{Sat}[\tau]$ and $\text{FO}[\tau] \setminus \text{Sat}[\tau]$ until one of them yields φ . □

Now we have the negative result concerning the proof calculus: If there would be a complete proof calculus for first-order logic in the finite, we could effectively enumerate all possible formal proofs and hence, also the sentences valid in the finite, contradicting Theorem 7.2.2.

7.2.2 Structures as Inputs

We fix the conventions which allow to regard finite structures as inputs to Turing machines. Note that, in general, structures are abstract objects and the same holds for their elements, so that there is no canonical way of representing structures by strings (or by sequences of strings). If the same structure


has different representations, the machine should produce the same answer on all the representations. We postpone this problem and first make life easier by restricting our attention to ordered structures. Recall the definition (cf. 1.A2):

Definition 7.2.3 Let $\{<\} \subseteq \tau_0 \subseteq \{<, S, \min, \max\}$ and $\tau_0 \subseteq \tau$. A τ -structure \mathcal{A} is *ordered* if the reduct $\mathcal{A}|_{\tau_0}$ is an ordering (that is, $<^{\mathcal{A}}$ is an ordering and S, \min , and \max , if present, are interpreted by the successor relation, the least and the last element of the ordering, respectively). $\mathcal{O}[\tau]$ is the class of ordered τ -structures. If ψ is a sentence in the vocabulary τ , $\text{ordMod}(\psi)$ denotes the class of ordered models of ψ , or equivalently, $\text{ordMod}(\psi) = \text{Mod}(\psi \wedge \psi_0)$, where ψ_0 is the conjunction of the ordering axioms for the vocabulary τ_0 . \square

Let $\mathcal{A} \in \mathcal{O}[\tau]$ be an ordered structure with $\|\mathcal{A}\| = n$. By passing to an isomorphic copy we can assume – and always will tacitly do so – that $A = \{0, \dots, n - 1\}$ and that $<^{\mathcal{A}}$ is the natural ordering on this set; that is, we identify or “label” the least element of $<^{\mathcal{A}}$ in \mathcal{A} with 0, its successor with 1, etc.³

Suppose $\tau = \tau_0 \dot{\cup} \tau_1$, say $\tau_1 = \{R_1, \dots, R_k, c_1, \dots, c_l\}$ (when writing τ_1 in this way we tacitly assume that the symbols in τ_1 are given in the order $R_1, \dots, R_k, c_1, \dots, c_l$), and τ_0 as in the preceding definition.

A *Turing machine for τ -structures* will have $1 + k + l$ *input tapes* and m *work tapes* for some $m \geq 1$. All tapes are bounded to the left and unbounded to the right. Their cells are numbered as indicated by⁴



The “virtual” cell is numbered by -1 and always contains the virtual letter α . All input tapes will contain an input word followed by the virtual letter ω indicating the end of the input word. Each tape has its own head. The heads can move independently of each other. Those on input tapes are *read-only* heads (i.e., they do not change the scanned letter), while those on the work tapes are *read-and-write* heads. The alphabet only contains the symbol “1”. Moreover, we identify “0” with “blank”; thus during a computation almost all cells contain 0.

³ By this convention, for isomorphic $\mathcal{A}, \mathcal{B} \in \mathcal{O}[\tau]$ a Turing machine M started with \mathcal{A} will have the same input inscriptions as M started with \mathcal{B} .

⁴ For the purposes of the next sections this numbering is more appropriate than that of the section on Trahtenbrot’s Theorem, where we numbered the virtual cell by 0.

With an ordered τ -structure \mathcal{A} we associate the following input inscriptions on the $1 + k + l$ input tapes (numbered from 0 to $k + l$): The 0-th tape, the “universe tape”, contains a sequence of 1’s of length $n := \|A\|$,

α	1	1	· · ·	1	ω
-1	0	1		$n-1$	n

For $1 \leq i \leq k$, the i -th input tape contains the information about $R := R_i$ coded as follows: Suppose R is r -ary, that is, $R^A \subseteq \{0, \dots, n-1\}^r$. Of course, $\|\{0, \dots, n-1\}^r\| = n^r$. Then, for $j < n^r$, the j -th cell will contain “1” just in case the j -th r -tuple in the lexicographic ordering of $\{0, \dots, n-1\}^r$ is in R . More formally: For $j < n^r$, let $|j|_r$ be the j -th r -tuple in the lexicographic ordering of $\{0, \dots, n-1\}^r$, i.e., look at the unique n -adic representation of j ,

$$j = j_1 \cdot n^{r-1} + j_2 \cdot n^{r-2} + \dots + j_{r-1} \cdot n + j_r \quad \text{with } 0 \leq j_i < n,$$

and set $|j|_r := (j_1, \dots, j_r)$. (If n is not clear from the context, we write $|j|_r^n$ instead of $|j|_r$.) Then the i -th input tape has the inscription

α	a_0	a_1	a_2	a_3	· · ·	a_{n^r-1}	ω
-1	0	1	2	3		n^r-1	n^r

where

$$a_j = 1 \quad \text{iff} \quad R^A |j|_r$$

(and hence, $a_j = 0$ iff not $R^A |j|_r$).

For $1 \leq i \leq l$, the $(k + i)$ -th input tape contains the binary representation of $j := c_i^A$ without leading zeros.

We say that a Turing machine M is *started with* \mathcal{A} , if the input tapes contain the information on \mathcal{A} in the way just described, the work tapes are empty, and each head scans the cell numbered 0 of its tape. As in the case of one-tape machines, M has finite sets $\text{State}(M)$ of states and $\text{Instr}(M)$ of instructions. $\text{State}(M)$ contains an *initial* (or *starting*) state s_0 , an *accepting* state s_+ , and a *rejecting* state s_- . *Instructions* now have the form

$$(*) \quad s b_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s' c'_1 \dots c'_m h_0 \dots h_{k+l+m}$$

with the meaning

If you are in state s , your heads scan $b_0; \dots, b_{k+l}$ on the input tapes and c_1, \dots, c_m on the work tapes, replace c_1, \dots, c_m by c'_1, \dots, c'_m , move the i -th head according to h_i and, finally, change to state s' .

Here, $s, s' \in \text{State}(M)$, $b_0, \dots, b_{k+l} \in \{0, 1, \alpha, \omega\}$, $c_1, \dots, c_m, c'_1, \dots, c'_m \in \{0, 1, \alpha\}$, and $h_0, \dots, h_{k+l+m} \in \{-1, 0, 1\}$. Moreover,

- if $b_j = \alpha$ then $h_j \neq -1$
(if the head scans the leftmost square it cannot move left)
- if $b_j = \omega$ then $h_j \neq 1$
- if $c_j = \alpha$ then $h_{k+l+j} \neq -1$ and $c'_j = \alpha$
- if $c_j \in \{0, 1\}$ then $c'_j \in \{0, 1\}$
- $s \neq s_+$ and $s \neq s_-$.

The *base* of the instruction in (*) is given by

$$s b_0 \dots b_{k+l} c_1 \dots c_m.$$

M is said to be *deterministic*, if no two distinct instructions in $\text{Instr}(M)$ have the same base. Sometimes, to emphasize that we do not require a machine to be deterministic, we speak of a *nondeterministic* machine.

The notions of accepting run, rejecting run, and of “ M accepts \mathcal{A} ” are adapted from section 7.2 in the obvious way. For a function $f : \mathbb{N} \rightarrow \mathbb{N}$ we say that M is *f time-bounded*, if for any \mathcal{A} accepted by M there is an accepting run of M , started with \mathcal{A} , of length at most $f(\|\mathcal{A}\|)$. And M is *f space-bounded*, if for all \mathcal{A} accepted by M there is an accepting run which uses at most $f(\|\mathcal{A}\|)$ squares on each work tape before stopping.

Let K be a class of ordered τ -structures. M *accepts* K if M accepts exactly those ordered τ -structures that lie in K . For classes of structures the definitions of PTIME (“polynomial time”), NPTIME (“nondeterministic polynomial time”), PSPACE (“polynomial space”) are introduced in the obvious way. For example,

- K is in PTIME iff there is a deterministic machine M and a polynomial $p \in \mathbb{N}[x]$ such that M accepts K and M is p time-bounded.

And we define

- K is in NLOGSPACE, “nondeterministic logarithmic space” (LOGSPACE, “deterministic logarithmic space”) iff there is a (deterministic) machine M and $d \geq 1$ such that M accepts K and is $d \cdot \log$ space-bounded ($\log n$ stands for the least natural number $\geq \log_2 n$).

Exercise 7.2.4 Denote by $l(\mathcal{A})$ the sum of the lengths of the input words on the input tapes of a machine started with \mathcal{A} . Show that in the definitions of the complexity classes mentioned above we can replace $\|\mathcal{A}\|$ by $l(\mathcal{A})$. \square

At the end of section 7.1 we have observed that for a class K of ordered structures and $m \geq 1$ the class

$$K_m := \{A \in K \mid \|A\| \geq m\}$$

is axiomatizable in a logic \mathcal{L} iff K is axiomatizable in \mathcal{L} . Analogously, for any of the complexity classes \mathcal{C} introduced so far, we have

$$K \in \mathcal{C} \quad \text{iff} \quad K_m \in \mathcal{C}$$

(since we can change a machine, without essentially effecting its time and space bounds, in such a way that it runs on a given finite set of inputs in a prescribed form; cf. Exercise 7.2.5 below). We give an application of this fact. Suppose that K is in PSPACE. Then K_2 is in PSPACE, too, say, it is accepted by a machine M that is q space-bounded, where $q(x) = a_s x^s + a_{s-1} x^{s-1} + \dots + a_1 x + a_0$. For suitable d , $q(n) \leq n^d$ for all $n \geq 2$. Thus M is x^d space-bounded. These observations will allow us to restrict ourselves to monic polynomials $p(x) = x^d$ when considering PSPACE (or PTIME, NPTIME).

Exercise 7.2.5 For fixed τ and $m \geq 1$ there is $c_m \in \mathbb{N}$ such that for any $K \subseteq \mathcal{O}[\tau]$ and any $f : \mathbb{N} \rightarrow \mathbb{N}$:

- (a) If K is accepted by an f time-bounded machine, then K_m is accepted by an $f + 2 \cdot (m + 1)$ time-bounded machine, where $f + 2 \cdot (m + 1)$ denotes the function whose value for n is $f(n) + 2 \cdot (m + 1)$.
- (b) If K_m is accepted by an f time-bounded machine, then K is accepted by an $f + c_m$ time-bounded machine.
- (c) K is accepted by an f space-bounded machine iff K_m is accepted by an f space-bounded machine. □

7.3 Logical Descriptions of Computations

Let K be a class of ordered τ -structures, $K \subseteq \mathcal{O}[\tau]$. We write $K \in \text{IFP}$ if K is axiomatizable in $\text{FO}(\text{IFP})$, and use similar notations for the other logics. Our main goal is to show

$K \in \text{LOGSPACE}$	iff	$K \in \text{DTC}$
$K \in \text{NLOGSPACE}$	iff	$K \in \text{TC}$
$K \in \text{PTIME}$	iff	$K \in \text{IFP}$
$K \in \text{NPTIME}$	iff	$K \in \Sigma_1^1$
$K \in \text{PSPACE}$	iff	$K \in \text{PFP}$.

(Σ_1^1 denotes the fragment of second-order logic consisting of the sentences of the form $\exists X_1 \dots \exists X_m \psi$ where ψ is first-order, cf. section 3.1).

These results provide the bridge between logic and complexity theory we have aimed at, by relating purely machine-oriented characterizations and

characterizations by means of logical definability. Section 7.5 exemplifies the methodological possibilities they offer.

In this section we prove the implications from left to right; the other implications will be settled in the next section.

Let \mathcal{C} be one of the complexity classes above and \mathcal{L} the logic associated to \mathcal{C} by the corresponding equivalence. Assume that $K \in \mathcal{C}$, and let M be a Turing machine witnessing that $K \in \mathcal{C}$. We are going to describe the behaviour of M by a formula φ_M of \mathcal{L} in such a way that for any ordered structure \mathcal{A} ,

$$\mathcal{A} \models \varphi_M \quad \text{iff} \quad M \text{ accepts } \mathcal{A}$$

and hence,

$$K = \text{ordMod}(\varphi_M).$$

The reader familiar with classical recursion theory will observe the analogy of the proofs given below with usual proofs showing the μ -recursiveness of Turing computable functions. There, the transition from one configuration to the next is coded by “simple” (e.g. primitive recursive) functions, and the μ -operator is used to detect the end of a computation. Similarly here, using the ordering relation, we can describe the transition from one configuration to the next by means of a “simple” logic, mostly first-order logic. The additional expressive power, for example the operator IFP, is used to find out whether the computation stops and to get its outcome. In recursion theory the proof of the equivalence shows that one application of the μ -operator suffices. Similarly we shall obtain, say, that every formula of FO(IFP) is equivalent to a formula with only one occurrence of IFP.

For the following we fix a vocabulary $\tau = \tau_0 \dot{\cup} \tau_1$ where, for simplicity, we assume that $\tau_0 = \{<, S, \min, \max\}$ and τ_1 is relational, $\tau_1 = \{R_1, \dots, R_k\}$ with r_i -ary R_i . For convenience we set $r_0 = 1$.

A Turing machine M for τ -structures has $1 + k$ input tapes and a certain number m of work tapes. In order to describe computations that M performs when started with a structure, we introduce *configurations* which contain the relevant data of possible situations of a computation. These data are:

- the current state
- the current inscriptions of the work tapes
- the current position of the heads on both the input and the work tapes.

An *accepting configuration* is a configuration with state s_+ . A configuration CONF' is said to be a *successor* of the configuration CONF , if an instruction of M allows M to pass from CONF to CONF' in one step. For convenience, we extend this definition by saying that any *accepting* configuration is a successor of itself. If M is deterministic then every configuration has at most one successor.

Let M be a (nondeterministic) Turing machine for τ -structures which is x^d space-bounded, that is, if M accepts an ordered structure \mathcal{A} , then there

is an accepting run that scans at most n^d squares on each work tape, where $n := \|A\|$. We may assume that $r_i \leq d$ for $i = 1, \dots, k$ (r_i being the arity of R_i). Fix a structure \mathcal{A} . By the remarks at the end of the last two sections when proving that a class of structures is axiomatizable in a logic or acceptable by a Turing machine of a certain complexity bound, we can restrict ourselves to sufficiently large finite structures. Here we look at structures \mathcal{A} such that for $n := \|A\|$ we have $n > k+m$ and $n > \|\text{State}(M)\|$. We assume that $\text{State}(M)$, the set of states of M , is an initial segment of the natural numbers and that $s_0 = 0$ is the starting state.

Let CONF be a configuration, where at most the n^d first cells of each work tape are not empty and where the heads scan one of these cells. A first attempt to code the contents of these cells could consist in dividing the relevant part of each work tape into $\frac{n^d}{\log n} =: r$ blocks of length $\log n$ and reading each block as a natural number $< n$ in binary representation. This would require variables x_1, \dots, x_r for each tape. Then a formula bearing the information on successive configurations would contain at least the variables x_1, \dots, x_r and would, therefore, depend on the cardinality n of the universe. We overcome this difficulty for PTIME, NPTIME, and PSPACE by using relation variables instead of individual variables.

To code the data of CONF we first introduce the “state relation” ST^{CONF} , the “end-of-tape relations” E_j^{CONF} , the “head relations” H_j^{CONF} , and the “inscription relations” I_j^{CONF} . The last ones are only introduced for the work tapes, since the inscriptions of input tapes are given by the input structure and kept fixed during the whole computation. We introduce: The unary relation

$$\text{ST}^{\text{CONF}} := \{s\} \text{ where } s \text{ is the state of CONF;}$$

for $0 \leq j \leq k+m$, the unary

$$\text{E}_j^{\text{CONF}} := \begin{cases} \{0\}, & \text{if the } j\text{-th head faces } \alpha \\ \{n-1\}, & \text{if the } j\text{-th head faces } \omega \\ \emptyset, & \text{otherwise;} \end{cases}$$

for $0 \leq j \leq k$, the r_j -ary⁵

$$\text{H}_j^{\text{CONF}} := \{|e|_{r_j} \mid 0 \leq e, \text{ the } j\text{-th head scans the } e\text{-th square and this does not contain } \omega\};$$

for $k+1 \leq j \leq k+m$, the d -ary

$$\text{H}_j^{\text{CONF}} := \{|e|_d \mid 0 \leq e, \text{ the } j\text{-th head scans the } e\text{-th square}\};$$

for $k+1 \leq j \leq k+m$, the d -ary

⁵ Recall that $|e|_{r_j}$ is the r_j -tuple given by the n -adic representation of e and that the square containing α is numbered -1 .

$$I_j^{\text{CONF}} := \{ |e|_d \mid 0 \leq e < n^d \text{ and the } e\text{-th square of the } j\text{-th work tape contains the symbol } 1 \}.$$

Obviously, CONF is uniquely determined by these relations. The starting configuration CONF_0 , for example, is given by

$$\text{ST}^{\text{CONF}_0} = \{0\}, \quad E_j^{\text{CONF}_0} = \emptyset, \quad H_j^{\text{CONF}_0} = \{(0, \dots, 0)\}, \quad \text{and } I_j^{\text{CONF}_0} = \emptyset.$$

For technical convenience we encode CONF in a single $(d + 2)$ -ary relation $C^{\text{CONF}} \subseteq \{0, \dots, n-1\}^{d+2}$ by joining the preceding relations as follows (we add two first coordinates to distinguish these relations and fill up with zeroes in the middle to get arity $(d + 2)$):⁶

$$\begin{aligned} C^{\text{CONF}} &:= \{(0, 0)\} \times \{\tilde{0}\} \times \text{ST}^{\text{CONF}} \\ &\cup \bigcup_{0 \leq j \leq k+m} \{(1, j)\} \times \{\tilde{0}\} \times E_j^{\text{CONF}} \\ &\cup \bigcup_{0 \leq j \leq k+m} \{(2, j)\} \times \{\tilde{0}\} \times H_j^{\text{CONF}} \\ &\cup \bigcup_{k+1 \leq j \leq k+m} \{(3, j)\} \times I_j^{\text{CONF}}. \end{aligned}$$

Clearly, given $C \subseteq \{0, \dots, n-1\}^{d+2}$, we can easily decide whether there is a configuration CONF of M , where only the first n^d cells of each work tape are relevant, such that $C = C^{\text{CONF}}$, or as we shortly say, whether C is an n^d -bounded configuration.

The following lemma now states how the behaviour of M can be described by suitable formulas.

Lemma 7.3.1 *Let M be a Turing machine which is x^d space-bounded. There is a first-order formula $\varphi_{\text{start}}(\bar{x})$ and there are first-order formulas $\varphi_{\text{succ}}(\bar{x}, X)$ and $\psi_{\text{succ}}(X, Y)$ (more precisely, second-order formulas without second-order quantifiers) such that for all sufficiently large $A \in \mathcal{O}[\tau]$ and $\bar{a} \in A^{d+2}$ we have:*

- (a) “ $\varphi_{\text{start}}(\bar{x})$ describes the starting configuration”: If C_0 denotes the starting configuration of M started with A then

$$A \models \varphi_{\text{start}}[\bar{a}] \quad \text{iff} \quad \bar{a} \in C_0.$$

- (b) “ $\varphi_{\text{succ}}(\bar{x}, X)$ describes the successor of X ”: If M is deterministic and C is an n^d -bounded configuration of M (where $n := \|A\|$) then

$$A \models \varphi_{\text{succ}}[\bar{a}, C] \quad \text{iff} \quad C \text{ has an } n^d\text{-bounded successor } C' \text{ and } \bar{a} \in C'.$$

⁶ By $\tilde{0}$ we denote constant sequences $0 \dots 0$ of appropriate length; in general, \tilde{c} will denote a sequence of the form $c \dots c$.

- (c) “ $\psi_{\text{succ}}(X, Y)$ expresses that Y is a successor of X ”: If C_1 is an n^d -bounded configuration of M and C_2 a further $(d + 2)$ -ary relation on A then

$$A \models \psi_{\text{succ}}[C_1, C_2] \text{ iff } C_2 \text{ is an } n^d\text{-bounded configuration of } M \text{ which is a successor of } C_1.$$

We postpone the proof of this lemma and first draw some consequences.

Theorem 7.3.2 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures. If K is in PSPACE then K is axiomatizable in FO(PFP).*

Proof. Let M be a deterministic machine witnessing $K \in \text{PSPACE}$. By previous remarks we can assume that M is x^d space-bounded for suitable d . We set

$$\varphi(\bar{x}, X) := (\neg \exists \bar{y} X \bar{y} \wedge \varphi_{\text{start}}(\bar{x})) \vee (\exists \bar{y} X \bar{y} \wedge \varphi_{\text{succ}}(\bar{x}, X)),$$

where φ_{start} and φ_{succ} are the formulas assigned to M in the preceding lemma. Let \mathcal{A} be an ordered structure and $n := \|\mathcal{A}\|$. By this lemma, $F_0^\varphi, F_1^\varphi, F_2^\varphi, \dots$ is the sequence $\emptyset, C_0, C_1, \dots$ where

- C_0 is the starting configuration
- if C_i is an n^d -bounded configuration of M with an n^d -bounded successor configuration C then $C_{i+1} = C$. In particular, if C_i is accepting then $C_i = C_{i+1} = C_{i+2} \dots$
- if C_i is an n^d -bounded configuration without a successor configuration or with a successor configuration which is not n^d -bounded, then $C_{i+1} = \emptyset, C_{i+2} = C_0, C_{i+3} = C_1, \dots$, that is, the sequence has no fixed-point.

Summarizing, we have

$$\begin{aligned} M \text{ accepts } \mathcal{A} & \text{ iff } F_\infty^\varphi \text{ is an accepting configuration} \\ & \text{ iff } F_\infty^\varphi \text{ is a configuration with state } s_+. \end{aligned}$$

“ F_∞^φ is a configuration with state s_+ ” is expressed by the formula (recall the $\text{ST}^{\text{CONF}}_{\text{-part}}$ of C^{CONF})

$$\exists y(\text{“}y \text{ is the } s_+\text{-th element of } <\text{”} \wedge [\text{PFP}_{\bar{x}, X} \varphi] \min \min \widetilde{\min} y).^7$$

We abbreviate it by $[\text{PFP}_{\bar{x}, X} \varphi] \min \min \widetilde{\min} s_+$. Then,

$$\begin{aligned} \mathcal{A} \in K & \text{ iff } M \text{ accepts } \mathcal{A} \\ & \text{ iff } \mathcal{A} \models [\text{PFP}_{\bar{x}, X} \varphi] \min \min \widetilde{\min} s_+, \end{aligned}$$

that is, $K = \text{ordMod}([\text{PFP}_{\bar{x}, X} \varphi] \min \min \widetilde{\min} s_+)$. □

⁷ Compare footnote 6.

Corollary 7.3.3 *Let $K \subseteq \mathcal{O}[\tau]$ be in PSPACE. Then K is axiomatizable by a sentence of FO(PFP) with only one occurrence of PFP. \square*

We now turn to PTIME and will realize that we can do better. Whereas accepting runs of polynomially space-bounded machines may have exponential length, we here only have to consider runs of polynomial length. This enables us to code the whole run into a single relation which can be obtained as the fixed-point of an inflationary process. On the logical side we therefore can replace PFP by IFP.

Consider a (finite or infinite) run C_0, C_1, \dots of an x^d time-bounded (and hence, x^d space-bounded) deterministic machine started with a structure of cardinality n . If the run accepts the structure, C_{n^d-1} must be an accepting configuration. The inflationary process indicated above is given by a formula $\varphi(\bar{v}, \bar{x}, Z)$ with

$$F_i^{(Z\bar{v}\bar{x} \vee \varphi)} = \bigcup_{\substack{m < i \\ C_m \text{ defined}}} \{|m|_d\} \times C_m, \quad F_\infty^{(Z\bar{v}\bar{x} \vee \varphi)} = \bigcup_{\substack{m < n^d \\ C_m \text{ defined}}} \{|m|_d\} \times C_m,$$

that is, we use the first d coordinates as time stamps when coding the run in one relation (as above, $|m|_d$ denotes the m -th tuple of $\{0, \dots, n-1\}^d$ in the lexicographic ordering).

Theorem 7.3.4 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures. If K is in PTIME then K is axiomatizable in FO(IFP).*

Proof. Let M be a deterministic machine witnessing $K \in \text{PTIME}$. We can assume that, for suitable d , M is x^d time-bounded. For $\bar{v} = v_0 \dots v_{d-1}$ we set

$$\varphi(\bar{v}, \bar{x}, Z) := (\bar{v} = \widetilde{\min} \wedge \varphi_{\text{start}}(\bar{x})) \vee \exists \bar{u} (S^d \bar{u} \bar{v} \wedge \varphi_{\text{succ}}(\bar{x}, Z\bar{u})).$$

Here, $\bar{v} = \widetilde{\min}$ abbreviates $v_0 = \min \wedge \dots \wedge v_{d-1} = \min$; $S^d \bar{u} \bar{v}$ stands for “ \bar{v} is the successor of \bar{u} in the lexicographic ordering”, and $\varphi_{\text{succ}}(\bar{x}, Z\bar{u})$ is obtained from $\varphi_{\text{succ}}(\bar{x}, X)$ by replacing subformulas $X\bar{t}$ by $Z\bar{u}\bar{t}$. Then we have for $A \in \mathcal{O}[\tau]$ with $n := \|A\|$:

$$\begin{aligned} A \in K & \text{ iff } M \text{ accepts } A \\ & \text{ iff the } (n^d-1)\text{-th configuration of } M, \text{ started} \\ & \quad \text{with } A, \text{ is defined and has state } s_+ \\ & \text{ iff } A \models [\text{IFP}_{\bar{v}\bar{x}, Z\varphi}] \widetilde{\max} \min \min \widetilde{\min} s_+, \end{aligned}$$

that is, K is the class of ordered models of a sentence of FO(IFP). \square

Corollary 7.3.5 *Let $K \subseteq \mathcal{O}[\tau]$ be in PTIME. Then K is axiomatizable in FO(IFP) by a sentence with only one occurrence of IFP. \square*

It is now easy to extend the previous result to NPTIME.

Theorem 7.3.6 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures. If K is in NPTIME then K is axiomatizable in SO by a Σ_1^1 -sentence.*

Proof. Choose M witnessing $K \in \text{NPTIME}$ and assume that M is x^d time-bounded. Then, for $\mathcal{A} \in \mathcal{O}[\tau]$ with $n := \|\mathcal{A}\|$,

- $\mathcal{A} \in K$ iff there is a run of M , started with \mathcal{A} , of length $\leq n^d$ that accepts \mathcal{A}
- iff there is a sequence C_0, \dots, C_{n^d-1} of n^d -bounded configurations of M , started with \mathcal{A} , such that C_0 is the starting configuration, C_{i+1} is a successor configuration of C_i , and s_+ is the state of C_{n^d-1}
- iff $\mathcal{A} \models \varphi$,

where φ is the sentence (the intended meaning of the second-order variable Z being $\bigcup_{m < n^d} \{|m|_d\} \times C_m$):

$$\varphi := \exists Z (\forall \bar{x} (Z \widetilde{\min} \bar{x} \leftrightarrow \varphi_{\text{start}}(\bar{x})) \wedge \forall \bar{u} \forall \bar{v} (S^d \bar{u} \bar{v} \rightarrow \psi_{\text{succ}}(Z \bar{u}_-, Z \bar{v}_-)) \wedge Z \widetilde{\max} \min \min \widetilde{\min} s_+).$$

□

We now come back to a proof of the key lemma 7.3.1.

Proof (of 7.3.1). Let M be an x^d space-bounded machine for τ -structures. Recall the specific coding of an n^d -bounded configuration CONF in a single relation C^{CONF} comprising relations $\text{ST}^{\text{CONF}}, \text{E}_j^{\text{CONF}}, \text{H}_j^{\text{CONF}}$, and I_j^{CONF} that contain the information on the state, the endmarks, the head positions, and the inscription of the work tapes. Let \bar{x} be the sequence of variables $xyx_1 \dots x_d$.

For (a) we can set

$$\begin{aligned} \varphi_{\text{start}}(\bar{x}) &:= \bar{x} = \tilde{0} && \text{“the state is } s_0\text{”} \\ &\vee (x = 2 \wedge 0 \leq y \leq k + m \wedge x_1 \dots x_d = \tilde{0}) && \text{“the heads scan the 0-th cell”} \end{aligned}$$

where, for example, $0 \leq y \leq k + m$ means that y is equal to or less than the $(k+m)$ -th element in the ordering $<$ and 0 stands for min. Also in the following we shall use such self-explanatory abbreviations.

For parts (b) and (c) we first introduce, for every instruction $\text{instr} \in \text{Instr}(M)$,

$$\text{instr} = sb_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s' c'_1 \dots c'_m h_0 \dots h_{k+l+m},$$

a formula $\varphi_{\text{instr}}(\bar{x}, X)$ which for n^d -bounded configurations X expresses

“ X has base $s\bar{b}\bar{c}$, and if the successor configuration according to instr is not n^d -bounded, then $\{\bar{x} \mid \varphi_{\text{instr}}(\bar{x}, X)\} = \emptyset$, otherwise $\{\bar{x} \mid \varphi_{\text{instr}}(\bar{x}, X)\}$ is this successor configuration”.

Moreover, we need a formula $\varphi_{\text{acc}}(X)$ which for n^d -bounded configurations X expresses

“ X is an accepting configuration”.

Using $\varphi_{\text{instr}}(\bar{x}, X)$ and $\varphi_{\text{acc}}(X)$, we get the desired formulas φ_{succ} and ψ_{succ} of parts (b) and (c) as (recall that we agreed to set $C_{m+1} = C_m$ for accepting configurations C_m)

$$\varphi_{\text{succ}}(\bar{x}, X) := (\varphi_{\text{acc}}(X) \wedge X\bar{x}) \vee \bigvee_{\text{instr} \in \text{instr}(M)} \varphi_{\text{instr}}(\bar{x}, X),$$

and

$$\psi_{\text{succ}}(X, Y) := (\varphi_{\text{acc}}(X) \wedge \forall \bar{x}(Y\bar{x} \leftrightarrow X\bar{x})) \vee \bigvee_{\text{instr} \in \text{instr}(M)} (\exists \bar{x} \varphi_{\text{instr}}(\bar{x}, X) \wedge \forall \bar{x}(Y\bar{x} \leftrightarrow \varphi_{\text{instr}}(\bar{x}, X))).$$

It remains to give $\varphi_{\text{instr}}(\bar{x}, X)$ and $\varphi_{\text{acc}}(X)$. We set

$$\varphi_{\text{acc}}(X) := X00\tilde{0}s_+.$$

The formula $\varphi_{\text{instr}}(\bar{x}, X)$ has the form

$$\varphi_{\text{instr}}(\bar{x}, X) := \varphi_{s, \bar{b}, \bar{c}}(X) \wedge \varphi_{s', \bar{c}', \bar{h}}(\bar{x}, X),$$

where, for an n^d -bounded configuration X , the formula $\varphi_{s, \bar{b}, \bar{c}}(X)$ expresses

“ X has base s, \bar{b}, \bar{c} ”,

and $\varphi_{s', \bar{c}', \bar{h}}(\bar{x}, X)$ expresses

“if the successor Y of X according to s', \bar{c}', \bar{h} is not n^d -bounded then $\{\bar{x} \mid \varphi_{s', \bar{c}', \bar{h}}(\bar{x}, X)\} = \emptyset$, else $\{\bar{x} \mid \varphi_{s', \bar{c}', \bar{h}}(\bar{x}, X)\} = Y$ ”.

For easier reading of the formulas below we introduce the following abbreviations:

ENDMARK $yz := X1y\tilde{0}z$	“the y -th head faces the endmark z ” (which is 0 for α and $n-1$ for ω)
HEAD $y\bar{z} := X2y\tilde{0}\bar{z}$	“the y -th head is on position $ \bar{z} $ ”
ONE $y\bar{z} := X3y\bar{z}$	“the y -th work tape contains 1 on position $ \bar{z} $ ”.

Now we take as $\varphi_{s, \bar{b}, \bar{c}}(X)$ the conjunction of

$X00\tilde{0}s$

“ s is the state”

$\bigwedge_{b_j=\alpha} \text{ENDMARK } j \text{ min} \wedge \bigwedge_{c_j=\alpha} \text{ENDMARK } (k+j) \text{ min}$
 “heads at the left end of a tape”

$\bigwedge_{b_j=\omega} \text{ENDMARK } j \text{ max}$
 “heads at the right end on input tapes”

$\bigwedge_{b_j=1} \exists x_1 \dots \exists x_{r_j} (\text{HEAD } j\tilde{0}x_1 \dots x_{r_j} \wedge R_j x_1 \dots x_{r_j})$
 “heads of input tapes facing a 1”

$\bigwedge_{b_j=0} \exists x_1 \dots \exists x_{r_j} (\text{HEAD } j\tilde{0}x_1 \dots x_{r_j} \wedge \neg R_j x_1 \dots x_{r_j})$
 “heads of input tapes facing a 0”

$\bigwedge_{c_j=1} \exists x_1 \dots \exists x_d (\text{HEAD } (k+j)x_1 \dots x_d \wedge \text{ONE } (k+j)x_1 \dots x_d)$
 “heads of work tapes facing a 1”

$\bigwedge_{c_j=0} \exists x_1 \dots \exists x_d (\text{HEAD } (k+j)x_1 \dots x_d \wedge \neg \text{ONE } (k+j)x_1 \dots x_d)$
 “heads of work tapes facing a 0”.

Finally, we take as $\varphi_{s', \bar{c}', \bar{h}}(\bar{x}, X)$ the conjunction $(\varphi_1 \wedge \varphi_2)$, where φ_1 is

$\bigwedge_{\substack{h_j=1 \\ k+1 \leq j \leq k+m}} \neg \text{HEAD } j \widetilde{\text{max}}$
 “heads of work tapes moving to the right do not face the $(n^d - 1)$ -th square”

and where φ_2 is the disjunction of

$(x = y = 0 \wedge x_1 \dots x_{d-1} = \tilde{0} \wedge x_d = s')$

“ s' is the new state”

$\bigvee_{k+1 \leq j \leq k+m} (\neg \text{HEAD } j x_1 \dots x_d \wedge \text{ONE } j x_1 \dots x_d \wedge x = 3 \wedge y = j)$
 “work tape content unchanged on squares not scanned”

$\bigvee_{c'_j=1} (\text{HEAD } j x_1 \dots x_d \wedge x = 3 \wedge y = k + j)$
 “new content 1 on scanned squares of a work tape”

$\bigvee_{h_j=1} (\text{ENDMARK } j\tilde{0} \wedge x = 2 \wedge y = j \wedge x_1 \dots x_d = \tilde{0})$
 “heads scanning α and moving to the right come to position 0”

$\bigvee_{h_j=0} (\text{ENDMARK } j\tilde{0} \wedge x = 1 \wedge y = j \wedge x_1 \dots x_d = \tilde{0})$
 “unchanged position of heads facing α ”

$$\begin{aligned}
 & \bigvee_{h_j=-1} (\text{ENDMARK } j \text{ max } \wedge x = 2 \wedge y = j \wedge x_1 \dots x_{d-r_j} = \tilde{0} \\
 & \quad \wedge x_{d-r_j+1} \dots x_d = \widetilde{\text{max}}) \\
 & \quad \text{“heads scanning } \omega \text{ and moving to the left come to position } n^{r_j}-1\text{”} \\
 & \bigvee_{h_j=0} (\text{ENDMARK } j \text{ max } \wedge x = 1 \wedge y = j \wedge x_1 \dots x_{d-1} = \tilde{0} \wedge x_d = \text{max}) \\
 & \quad \text{“unchanged position of heads facing } \omega\text{”} \\
 & \bigvee_{h_j=-1} (\text{HEAD } j\tilde{0} \wedge x = 1 \wedge y = j \wedge x_1 \dots x_d = \tilde{0}) \\
 & \quad \text{“heads scanning } \alpha \text{ from their new position”} \\
 & \bigvee_{\substack{j \leq k \\ h_j=1}} (\text{HEAD } j\tilde{0} \underbrace{\text{max} \dots \text{max}}_{r_j\text{-times}} \wedge x = 1 \wedge y = j \wedge x_1 \dots x_{d-1} = \tilde{0} \\
 & \quad \wedge x_d = \text{max}) \\
 & \quad \text{“heads of input tapes scanning } \omega \text{ from their new position”} \\
 & \bigvee_{j \leq k+m} \exists u_1 \dots \exists u_d (“x_1 \dots x_d = u_1 \dots u_d + h_j” \wedge \text{HEAD } ju_1 \dots u_d \wedge x = 2 \\
 & \quad \wedge y = j) \\
 & \quad \text{“new head position of heads on ‘interior’ squares.”} \quad \square
 \end{aligned}$$

We come to the corresponding results for LOGSPACE and NLOGSPACE. Recall that $K \in \text{NLOGSPACE}$ means that there is a (nondeterministic) machine M and some $d \geq 1$ such that M accepts K and is $d \cdot \log$ space-bounded. Every natural number $i < n$ codes a word over $\{0, 1\}$ of length $\log n$, namely its binary representation $|i|_{\log n}^2$ of length $\log n$. Thus, using d variables, we can represent the relevant contents of a work tape. Moreover, by restricting ourselves to sufficiently large structures \mathcal{A} , we can assume that $d \cdot \log n < n$ (where $n := \|A\|$). Hence, each head position can be represented by a single number $< n$. Altogether, we can describe the data of a configuration by a sequence of natural numbers $< n$ of length independent of n , where we agree to use the first number to represent the state. The exact definitions will be given later when proving the following lemma.

Lemma 7.3.7 *Let M be $d \cdot \log$ space-bounded. Then there are formulas $\chi_{\text{start}}(\bar{x})$ of FO and $\chi_{\text{succ}}(\bar{x}, \bar{x}')$ of FO(DTC) such that for all sufficiently large $\mathcal{A} \in \mathcal{O}[\tau]$ and \bar{a} in A ,*

- (a) $\mathcal{A} \models \chi_{\text{start}}[\bar{a}]$ iff \bar{a} is the (description of the) starting configuration.
- (b) For any $(d \cdot \log \|A\|)$ -bounded configuration \bar{a} and any \bar{b} ,

$$\mathcal{A} \models \chi_{\text{succ}}[\bar{a}, \bar{b}] \quad \text{iff } \bar{b} \text{ is a } (d \cdot \log \|A\|)\text{-bounded successor configuration of } \bar{a}.$$

Before we give a proof we derive some consequences.

Theorem 7.3.8 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures. If $K \in \text{LOGSPACE}$ then K is axiomatizable in FO(DTC).*

Proof. Let M be a deterministic machine witnessing $K \in \text{LOGSPACE}$, say, M is $d \log$ space-bounded. Let χ_{start} and χ_{succ} be the formulas corresponding to M according to the preceding lemma. Then, by (a) and (b) of this lemma, we have for $\mathcal{A} \in \mathcal{O}[\tau]$:

M accepts \mathcal{A}

- iff there is a sequence $\bar{a}_0, \dots, \bar{a}_k$ of $(d \cdot \log \|A\|)$ -bounded configurations such that \bar{a}_0 is the starting configuration, \bar{a}_{i+1} is the successor configuration of \bar{a}_i , and \bar{a}_k is an accepting configuration
- iff $\mathcal{A} \models \exists \bar{x} (\chi_{\text{start}}(\bar{x}) \wedge \exists \bar{x}' ([\text{DTC}_{\bar{x}, \bar{x}'} \chi_{\text{succ}}(\bar{x}, \bar{x}')] \bar{x} \bar{x}' \wedge x'_1 = s_+))$.

Hence, K is the class of ordered models of a sentence of $\text{FO}(\text{DTC})$. \square

Theorem 7.3.9 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures. If $K \in \text{NLOGSPACE}$ then K is axiomatizable in $\text{FO}(\text{TC})$.*

Proof. Let M be a machine witnessing $K \in \text{NLOGSPACE}$. Since M is non-deterministic we just have to replace DTC by TC in the last proof. In fact, we have for $\mathcal{A} \in \mathcal{O}[\tau]$

M accepts \mathcal{A}

- iff $\mathcal{A} \models \exists \bar{x} (\chi_{\text{start}}(\bar{x}) \wedge \exists \bar{x}' ([\text{TC}_{\bar{x}, \bar{x}'} \chi_{\text{succ}}(\bar{x}, \bar{x}')] \bar{x} \bar{x}' \wedge x'_1 = s_+))$.

Since $\text{FO}(\text{DTC}) \leq \text{FO}(\text{TC})$, we have obtained a sentence of $\text{FO}(\text{TC})$ axiomatizing K . \square

Denote by $\text{FO}(\text{posTC})$ the class of formulas of $\text{FO}(\text{TC})$ which only contain positive occurrences of TC, that is, each such occurrence is in the scope of an even number of negation symbols. In 8.6.12 we will see that $\text{FO}(\text{DTC}) \leq \text{FO}(\text{posTC})$. Thus the preceding proof yields

Corollary 7.3.10 *If a class of ordered structures is in NLOGSPACE then it is axiomatizable by a sentence of $\text{FO}(\text{posTC})$.* \square

We turn to a proof of Lemma 7.3.7. As already remarked above, in the case of log space-bounded machines we code configurations not by relations but by numbers. This change requires to define some arithmetical functions and predicates in $\text{FO}(\text{DTC})$. The next two lemmas serve this purpose.

Lemma 7.3.11 *There are $\text{FO}(\text{DTC})$ -formulas*

$$\varphi_+(x, y, z), \varphi_-(x, y, z), \varphi_2(x, y), \text{ and } \varphi_{\log(\text{universe})}(x)$$

such that for any ordered structure \mathcal{A} with $A = \{0, \dots, \|A\| - 1\}$ and any $a, b, c \in A$,

$$\begin{array}{ll} \mathcal{A} \models \varphi_+[a, b, c] & \text{iff } a + b = c \\ \mathcal{A} \models \varphi_-[a, b, c] & \text{iff } a \cdot b = c \\ \mathcal{A} \models \varphi_2[a, b] & \text{iff } 2^a = b \\ \mathcal{A} \models \varphi_{\log(\text{universe})}[a] & \text{iff } a = \log \|A\|. \end{array}$$

Proof. For better readability, instead of describing the natural numbers in terms of the ordering, we use constants $1, 2, \dots$

Given numbers x and y , the path

$$(0, x) \rightarrow (1, x+1) \rightarrow (2, x+2) \rightarrow \dots \rightarrow (y, x+y)$$

from $(0, x)$ to $(y, x+y)$ shows that as $\varphi_+(x, y, z)$ we can take the formula

$$(y = \min \wedge z = x) \vee [\text{DTC}_{uv, u'v'} (Suu' \wedge Svv')] \min xyz.$$

Similarly, the path

$$(0, 0) \rightarrow (1, x) \rightarrow (2, 2 \cdot x) \rightarrow (3, 3 \cdot x) \rightarrow \dots \rightarrow (y, y \cdot x)$$

shows that we can set

$$\begin{aligned} \varphi_-(x, y, z) &:= (y = \min \wedge z = \min) \\ &\vee [\text{DTC}_{uv, u'v'} (Suu' \wedge \varphi_+(v, x, v'))] \min \min yz \end{aligned}$$

and the path

$$(0, 1) \rightarrow (1, 2) \rightarrow (2, 4) \rightarrow \dots \rightarrow (x, 2^x)$$

shows that we can set

$$\begin{aligned} \varphi_2(x, y) &:= (x = \min \wedge y = 1) \\ &\vee [\text{DTC}_{uv, u'v'} (Suu' \wedge \varphi_-(v, 2, v'))] \min 1xy. \end{aligned}$$

Finally, let

$$\varphi_{\log(\text{universe})}(x) := \neg \exists y \varphi_2(x, y) \wedge \forall z (z < x \rightarrow \exists y \varphi_2(z, y)). \quad \square$$

Let $l := \log n - 1$. Then $2^l < n$. Recall that for m with $m < 2^l$ and $m_0, \dots, m_{l-1} \in \{0, 1\}$ we have

$$|m|_l^2 = m_0 \dots m_{l-1} \quad \text{iff} \quad m = m_0 \cdot 2^{l-1} + m_1 \cdot 2^{l-2} + \dots + m_{l-2} \cdot 2 + m_{l-1}.$$

We then say that m_k is the k -th digit of $|m|_l^2$. If n , and hence l , is clear from the context, we denote $|m|_l^2$ by $[m]$. We write \bar{u} for $u_0 \dots u_d$ and $\bar{u} < 2^l$ for $u_0 < 2^l \wedge \dots \wedge u_d < 2^l$; similarly with \bar{u}', \bar{x} , and \bar{x}' .

Lemma 7.3.12 *There are formulas of FO(DTC) which, in ordered structures \mathcal{A} , define the following relations (where $n = \|\mathcal{A}\|$, $l = \log n - 1$):*

- One mk iff $m < 2^l$, $k < l$, and the k -th digit of $[m]$ is 1
- Zero mk iff $m < 2^l$, $k < l$, and the k -th digit of $[m]$ is 0
- One $_d \bar{u}k$ iff $\bar{u} < 2^l$, $k < (d+1) \cdot l$, and the k -th digit of the concatenation $[u_0] \dots [u_d]$ is 1
- Zero $_d \bar{u}k$ iff $\bar{u} < 2^l$, $k < (d+1) \cdot l$, and the k -th digit of the concatenation $[u_0] \dots [u_d]$ is 0
- Equal $_d \bar{u}k\bar{u}'$ iff $\bar{u}, \bar{u}' < 2^l$, $k < (d+1) \cdot l$, and the words $[u_0] \dots [u_d]$ and $[u'_0] \dots [u'_d]$ differ at most at the k -th position.

We denote the corresponding formulas by $\varphi_{\text{one}}(x, z)$, $\varphi_{\text{zero}}(x, z)$, $\varphi_{d\text{-one}}(\bar{x}, z)$, $\varphi_{d\text{-zero}}(\bar{x}, z)$, and $\varphi_{d\text{-equal}}(\bar{x}, z, \bar{x}')$, respectively.

Proof. Note that for $m < 2^l$ and $k < l$ the k -th digit of $[m]$ is 1 iff

$$\exists y \in \mathbb{N} \exists z \in \mathbb{N} (m = y \cdot 2^{l-k} + z \wedge 2^{l-k-1} \leq z < 2^{l-k}).$$

Using this fact one easily shows, for example, the following equivalences, which immediately can be formalized with the help of the formulas of the preceding lemma.

$$\begin{aligned} \text{One } mk & \quad \text{iff } m < 2^l \wedge k < l \wedge \\ & \quad \exists y \exists z (m = y \cdot 2^{l-k} + z \wedge 2^{l-k-1} \leq z < 2^{l-k}) \\ \text{One}_1 u_0 u_1 k & \quad \text{iff } (u_0 < 2^l \wedge u_1 < 2^l \wedge k < l \wedge \text{One } u_0 k) \vee \\ & \quad (l \leq k < 2 \cdot l \wedge \text{One } u_1 (k-l)) \\ \text{Equal}_d \bar{u} k \bar{u}' & \quad \text{iff } \bar{u} < 2^l \wedge \bar{u}' < 2^l \wedge k < (d+1) \cdot l \wedge \\ & \quad \forall i ((i < (d+1) \cdot l \wedge i \neq k) \rightarrow (\text{One}_d \bar{u} i \leftrightarrow \text{One}_d \bar{u}' i)). \quad \square \end{aligned}$$

Finally, we come to the

Proof (of 7.3.7). Let M be a logarithmically space-bounded machine for τ -structures, say, M is $d \cdot \log$ space-bounded. For simplicity we assume that $\tau = \{R\}$ with binary R and that M only has one work tape. We restrict ourselves to structures of cardinality n with $n > d \cdot \log n$, $(d+1) \cdot l \geq d \cdot \log n$, and $n > s_M + 1$ (recall that $l := \log n - 1$ and that $\{0, \dots, s_M\}$ is the set of states of M). When M is started with a structure \mathcal{A} , where $A = \{0, \dots, n-1\}$, we can code the data of a resulting configuration by a tuple

$$(z, u_\alpha, u_\omega, u, v_\alpha, v_\omega, v_0, v_1, w_\alpha, w, y_0, \dots, y_d)$$

where

- z is the state
- u_α, u_ω, u code the position of the head on the 0-th input tape (the “universe tape”) by

$$u_\alpha = \begin{cases} 0 & \text{if the head does not face } \alpha \\ n-1 & \text{if the head faces } \alpha \end{cases}$$

$$u_\omega = \begin{cases} 0 & \text{if the head does not face } \omega \\ n-1 & \text{if the head faces } \omega \end{cases}$$

and u is the number of the cell faced by the head if it is an interior one, otherwise $u = 0$

- similarly, $v_\alpha, v_\omega, v_0, v_1$ code the position of the head on the first input tape, the tape for the binary relation R ; the variables v_0, v_1 represent the head position $v_0 \cdot n + v_1$

- w_α, w code the position of the head on the work tape (note that $d \cdot \log n < n$)
- the concatenation $[y_0] \dots [y_d]$ is the inscription of the first $(d + 1) \cdot l$ cells of the work tape.

Sometimes, for notational simplicity, we denote the sequence $zu_\alpha u_\omega \dots y_d$ simply by \bar{x} .

Concerning part (a) of the lemma, we can set

$$\chi_{\text{start}}(\bar{x}) := \bar{x} = \bar{0}.$$

Concerning part (b), we define $\chi_{\text{succ}}(\bar{x}, \bar{x}')$ by

$$\chi_{\text{succ}}(\bar{x}, \bar{x}') := \chi_{\text{acc}}(\bar{x}, \bar{x}') \vee \bigvee_{\text{instr} \in \text{instr}(M)} \chi_{\text{instr}}(\bar{x}, \bar{x}')$$

where

$$\chi_{\text{acc}}(\bar{x}, \bar{x}') := (x_1 = s_+ \wedge \bar{x}' = \bar{x})$$

(which in case \bar{x} is a configuration expresses that \bar{x} is accepting and $\bar{x}' = \bar{x}$) and where for every instruction

$$\text{instr} = sb_0b_1c_1 \rightarrow s'c'_1h_0h_1h_2$$

$\chi_{\text{instr}}(\bar{x}, \bar{x}')$ is a formula which, in case \bar{x} is a $(d \cdot \log)$ -bounded configuration, expresses that \bar{x} has base $sb_0b_1c_1$, the successor configuration of \bar{x} according to instr is $(d \cdot \log)$ -bounded and is \bar{x}' . As an example, we explicitly give $\chi_{\text{instr}}(\bar{x}, \bar{x}')$ for

$$\text{instr} = s1\alpha 1 \rightarrow s'0(-1)11,$$

namely as the conjunction of

$$z = s$$

“ s is the state”

$$u_\alpha = \min \wedge u_\omega = \min$$

“the head of the 0-th input tape faces an interior cell”

$$v_\alpha = \max \wedge v_\omega = \min \wedge v_0 = \min \wedge v_1 = \min$$

“the head of the first input tape faces α ”

$$w_\alpha = \min \wedge \text{One}_d y_0 \dots y_d w$$

“the head of the work tape faces a 1”

$$z' = s'$$

“ s' is the new state”

$$u'_\omega = \min \wedge ((u > 0 \wedge Su'u \wedge u'_\alpha = \min) \vee (u = 0 \wedge u' = 0 \wedge u'_\alpha = \max))$$

“new head position of the 0-th input tape”

$$v'_\alpha = \min \wedge v'_0 = \min \wedge v'_1 = \min \wedge v'_\omega = \min$$

“new head position of the first input tape is cell 0”

$$w'_\alpha = \min \wedge Sww' \wedge w' < d \cdot \log n,$$

i.e. $w'_\alpha = \min \wedge Sww' \wedge \exists x(\varphi_{\log(\text{universe})}(x) \wedge w' < d \cdot x)$

“new head position of the work tape is within the bounds”

$$\text{Zero}_d y'_0 \dots y'_d w$$

“new content of cell scanned on the work tape”

$$\text{Equal}_d y_0 \dots y_d w y'_0 \dots y'_d$$

“work tape content unchanged on cells not scanned”. □

7.4 The Complexity of the Satisfaction Relation

Suppose, for example, that the class K of ordered structures is axiomatizable by the FO(IFP)-sentence φ , $K = \{\mathcal{A} \in \mathcal{O}[\tau] \mid \mathcal{A} \models \varphi\}$. We aim at showing that $K \in \text{PTIME}$, that is, for fixed φ we want to prove that the satisfaction relation $\mathcal{A} \models \varphi$ can be decided in time polynomially bounded in $\|\mathcal{A}\|$. One also says that φ has a polynomial time *model-checker*. To simplify the corresponding algorithms we start with some general remarks. Note that manipulations in algorithms of the kind we now describe do not destroy polynomial time and logarithmic space bounds.

- (1) Using an additional work tape W' it is possible at any time of a computation to move the head of a given work tape W to the rightmost square which the head of the tape has scanned so far. (In fact, change the given program so that the head on W' moves in the same way as the head on W , but always prints the symbol 1.)
- (2) By (1) it is possible at any time of a computation to erase the content of a work tape (note that the additional work tape used in (1) can be cleared in a trivial way); in particular, one can change a program – without changing the accepted class – such that all work tapes are empty whenever the program stops.
- (3) The content of a worktape W can be copied to an empty tape W_1 (using (1), bring the corresponding heads H and H_1 to the rightmost cell scanned by H and copy the content cell by cell).
- (4) In our applications the 0-th input tape has the inscription $\underbrace{1 \dots 1}_{n\text{-times}}$, where n is the cardinality of the structure we consider. One can write the binary representation of n (of length $\leq \log n$) on a work tape. We say “a counter is set to n ”. Similarly, a counter can be set to n^d for any fixed $d \geq 1$.

Let \mathcal{L} be one of the logics considered in the preceding section and \mathcal{C} the corresponding complexity class. We want to show that for any sentence of

\mathcal{L} the class K of its ordered models is in \mathcal{C} . We even show that there is a machine M *strongly witnessing* $K \in \mathcal{C}$, that is,

- M accepts K ;
- for any $\mathcal{A} \in \mathcal{O}[\tau]$ every run of M , started with \mathcal{A} , stops at s_+ or s_- ; in particular, if M is deterministic then M decides K ;
- for any $\mathcal{A} \in \mathcal{O}[\tau]$ every run of M satisfies the time or space bounds characteristic for \mathcal{C} .

The proof that the class of ordered models of a sentence φ of \mathcal{L} is in \mathcal{C} proceeds by induction on φ . So we have to deal with formulas, too. We therefore introduce the following notation: For a formula $\varphi(x_1, \dots, x_l, Y_1, \dots, Y_r)$ we let

$$\text{ordMod}(\varphi) := \{(\mathcal{A}, a_1, \dots, a_l, P_1, \dots, P_r) \mid \mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \varphi[\bar{a}, \bar{P}]\},$$

that is, we consider the ordered models of $\varphi(c_1, \dots, c_l, P_1, \dots, P_r)$, a sentence in an enlarged vocabulary.

Theorem 7.4.1 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures.*

- (a) *If $K \in \text{DTC}$ then $K \in \text{LOGSPACE}$.*
- (b) *If $K \in \text{posTC}$ then $K \in \text{NLOGSPACE}$.*

Proof. By induction on the corresponding formulas φ we show that the class of ordered models of φ is in LOGSPACE and NLOGSPACE, respectively, and that there exists a machine strongly witnessing this fact. We handle both cases simultaneously. By passing to an equivalent formula we can assume that in formulas of FO(posTC) the TC operation does not occur in the scope of any negation symbol (the new formula may also contain \wedge, \vee).

Suppose that φ is atomic, say for simplicity, $\varphi = Rxy$. We show that there is a machine M strongly witnessing that

$$\{(\mathcal{A}, i, j) \mid \mathcal{A} \in \mathcal{O}[\tau], R^A ij\} \in \text{LOGSPACE}.$$

Let $(\mathcal{A}, i, j) \in \mathcal{O}[\tau \cup \{c, d\}]$ with $A = \{0, 1, \dots, n-1\}$ be given. Note that the information, whether $R^A ij$ holds, is to be found in the $(i \cdot n + j)$ -th square of the input tape corresponding to R , and (the binary representations of) i and j are available on the input tapes corresponding to c and d . Now it should be clear how a machine strongly witnessing that $\text{ordMod}(Rxy) \in \text{LOGSPACE}$ can be designed.

$\varphi = \neg\psi$: By the remarks above, ψ does not contain TC and hence, is in FO(DTC). By induction hypothesis, there is a machine M strongly witnessing that $\text{ordMod}(\psi) \in \text{LOGSPACE}$. For φ just interchange the roles of s_+ and s_- in M .

$\varphi(x_1, \dots, x_l) = (\psi \vee \chi)$: By induction hypothesis, there are corresponding machines M_ψ for $\psi(x_1, \dots, x_l)$ and M_χ for $\chi(x_1, \dots, x_l)$. Let M be a machine that first carries out the computation of M_ψ and then that of M_χ accepting

the input in case at least one, M_ψ or M_χ , accepts, and rejecting otherwise. (After the computation of M_ψ the work tapes are erased as explained in remark (2) above.)

$\varphi = (\psi \wedge \chi)$: similarly.

$\varphi(x_1, \dots, x_l) = \exists x\psi$: By induction hypothesis, there is a corresponding machine M_0 for $\psi(x_1, \dots, x_l, x)$. A machine M for φ operates as follows: Suppose M is started with an ordered structure $(\mathcal{A}, a_1, \dots, a_l)$ where $A = \{0, \dots, n-1\}$. Then, for $i = 0, \dots, n-1$, M writes the binary representation of i on a work tape and checks, using M_0 , whether $\mathcal{A} \models \psi[a_1, \dots, a_l, i]$. (Note that the binary representation of i on the work tape does not carry an endmark ω as required on the corresponding input tape of M_0 . To remind the end of the representation of i , proceed as explained in remark (1) above.) If the answer is positive at least once, M stops in state s_+ , otherwise in s_- .

$\varphi = \forall x\psi$: similarly.

$\varphi = [\text{DTC}_{\bar{x}, \bar{y}} \psi] \bar{s} \bar{t}$, where ψ is a formula of $\text{FO}(\text{DTC})$: For simplicity, we assume that the free variables of ψ are among \bar{x}, \bar{y} and that $\bar{x} = x, \bar{y} = y, \bar{s} = s$, and $\bar{t} = t$. Choose a machine M_0 strongly witnessing that $\text{ordMod}(\psi) \in \text{LOGSPACE}$. Given \mathcal{A} with $A = \{0, \dots, n-1\}$, if there is a ψ -path from s to t , there is one of length $\leq n$. Therefore, the machine M , we aim at, can be organized as follows: It writes $i := s$ on a work tape and sets a counter to n , which is used to invoke a subroutine at most n times. M rejects in case the counter becomes negative. Using M_0 , the subroutine checks for $j = 0, \dots, n-1$ whether $\mathcal{A} \models \psi[i, j]$ holds for exactly one j ; if not, M rejects. Otherwise, M checks whether j equals t ; in the affirmative case M accepts. In the negative case, M sets $i := j$ and reduces the counter by one.

$\varphi = [\text{TC}_{\bar{x}, \bar{y}} \psi] \bar{s} \bar{t}$: Once more, for simplicity, we assume that the free variables of ψ are among \bar{x}, \bar{y} and that $\bar{x} = x, \bar{y} = y, \bar{s} = s$, and $\bar{t} = t$. Choose a machine M_0 strongly witnessing $\text{ordMod}(\psi) \in \text{NLOGSPACE}$. We give the basic idea underlying the construction of a machine M for φ : A counter is set to n (again, we assume that $A = \{0, 1, \dots, n-1\}$), which is used to carry out a subroutine at most n times. M will stop in the rejecting state in case the counter becomes negative. M writes $i := s$ on a work tape. The subroutine starts choosing an element $j \in \{0, \dots, n-1\}$ nondeterministically (by using a counter to randomly write a $\{0, 1\}$ word of length $\log n$ on a tape). Then it checks, using M_0 , whether $\psi(i, j)$ holds. If not, M stops in state s_- . Otherwise, if $j = t$, M stops in s_+ , and if $j \neq t$, it sets $i := j$. \square

Theorem 7.4.2 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures.*

- (a) *If $K \in \text{IFP}$ then $K \in \text{PTIME}$.*
- (b) *If $K \in \text{PFP}$ then $K \in \text{PSPACE}$.*

Proof. Again, we give the proof by induction on the formula φ axiomatizing the class K . The cases where φ is atomic, $\neg\psi$, $(\psi \vee \chi)$, or $\exists x\psi$ are handled as in the preceding proof. The corresponding machines are polynomially time-bounded or space-bounded if the machines used in the induction hypotheses

are. Now, for part (a), suppose that $\varphi = [\text{IFP}_{\bar{x}, X} \psi(\bar{x}, X)] \bar{t}$ where X is r -ary, and that M_0 is a machine strongly witnessing that

$$\{(\mathcal{A}, \bar{a}, R) \mid \mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \psi[\bar{a}, R]\} \in \text{PTIME}$$

(for simplicity, we assume that the free variables of ψ are among \bar{x}, X). The machine M , we look for, contains a subroutine that uses work tapes W and W' . If started with a word of length n^r on W – the code of an r -ary relation R – and empty W' , the subroutine writes, invoking the machine M_0 , the code of

$$R' := \{\bar{a} \mid \mathcal{A} \models (X\bar{x} \vee \psi)[\bar{a}, R]\}$$

on the tape W' without changing the content of W .

The machine M for φ operates as follows: It sets $R := \emptyset$ and uses the subroutine to calculate R' . If $R = R'$ it checks whether $R\bar{t}$ or not $R\bar{t}$ and accepts or rejects, respectively. Otherwise, it sets $R := R'$, erases the content of W' , and again starts the subroutine. Note that $R = R'$ will be achieved after at most n^d calls to the subroutine (by 7.1.1(c)).

For part (b) assume that $\varphi = [\text{PFP}_{\bar{x}, X} \psi(\bar{x}, X)] \bar{t}$ with r -ary X and that M_0 is a machine strongly witnessing $\{(\mathcal{A}, \bar{a}, R) \mid \mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \psi[\bar{a}, R]\} \in \text{PSPACE}$. Given a structure \mathcal{A} , the operation F^ψ associated with ψ satisfies $F_{2^{n^r}-1}^\psi = F_{2^{n^r}}^\psi$ (and this set is the fixed-point F_∞^ψ) or $F_\infty^\psi = \emptyset$ (cf. 7.1.1(b)). The machine M for φ starts its computation on \mathcal{A} by setting a counter to $2^{n^r} - 1$ (that is, it writes the word $1 \dots 1$ of length n^r on a work tape). Then it proceeds as in the IFP-case, but now using the counter to ensure that the subroutine which here evaluates

$$R' := \{\bar{a} \mid \mathcal{A} \models \psi[\bar{a}, R]\},$$

is invoked at most $2^{n^r} - 1$ times. When the counter gets negative, it checks whether $R = R'$ and whether $R\bar{t}$. If both questions are answered positively it accepts, otherwise it rejects. \square

Theorem 7.4.3 *Let $K \subseteq \mathcal{O}[\tau]$ be a class of ordered structures.*

- (a) *If $K \in \Sigma_1^1$ then $K \in \text{NPTIME}$.*
- (b) *If $K \in \text{SO}$ then $K \in \text{PSPACE}$.*

Proof. (a) Let $K = \text{Mod}(\varphi)$ where $\varphi = \exists X_1 \dots \exists X_l \psi$, ψ is first-order, and the arity of X_i is r_i . By 7.4.1(a), there is a machine M_0 strongly witnessing that $\text{Mod}(\psi(X_1, \dots, X_l))$ is in PTIME (even in LOGSPACE). The machine M for φ , started with $\mathcal{A} \in \mathcal{O}[\tau]$, nondeterministically writes words over $\{0,1\}$ of length n^{r_1}, \dots, n^{r_l} on different work tapes, which are intended as codes of interpretations P_1, \dots, P_l of X_1, \dots, X_l . Then, using M_0 , it checks whether $\mathcal{A} \models \psi[P_1, \dots, P_l]$ or not, and stops in an accepting or rejecting state, respectively.

(b) Let $K = \text{Mod}(\varphi)$ for the formula φ of SO. To gain a machine M witnessing $K \in \text{PSPACE}$ we proceed by induction on φ . For φ atomic or of the form $\neg\psi$, $(\psi \vee \chi)$, or $\exists x\psi$ we argue as in the proof of 7.4.1. For $\varphi = \exists X\psi$ with r -ary X the machine M writes the word $1 \dots 1$ of length n^r on a work tape W ; then it systematically decreases this word, checking in each case with a polynomially space-bounded machine for ψ , whether ψ holds, if the interpretation of X is given by the tape W . \square

7.5 The Main Theorem and Some Consequences

To summarize the results of the preceding sections we introduce the following notion:

Definition 7.5.1 A logic \mathcal{L} *captures* a complexity class \mathcal{C} if for all τ with $\langle \in \tau$ and $K \subseteq \mathcal{O}[\tau]$ we have

$$K \in \mathcal{C} \quad \text{iff} \quad K \text{ is axiomatizable in } \mathcal{L}. \quad \square$$

- Theorem 7.5.2 (Main Theorem)** (a) $\text{FO}(\text{DTC})$ *captures* LOGSPACE.
 (b) $\text{FO}(\text{TC})$ *captures* NLOGSPACE.
 (c) $\text{FO}(\text{IFP})$ *captures* PTIME.
 (d) Σ_1^1 *captures* NPTIME.
 (e) $\text{FO}(\text{PFP})$ *captures* PSPACE.

Note that we have proved the theorem up to part (b). We only have shown that

$$\text{FO}(\text{posTC}) \text{ captures NLOGSPACE.}$$

At the end of this section we will prove that

$$\text{FO}(\text{posTC}) \equiv \text{FO}(\text{TC}) \text{ on ordered structures.}$$

Then the proof will be finished.

The main theorem gives descriptive characterizations of some important complexity classes by certain extensions of first-order logic. Is there a well-known complexity class that, in this sense, corresponds to first-order logic itself? We have mentioned already several times that first-order logic has rather limited possibilities to speak about inductive or recursive procedures. In fact, it has been proved that first-order logic captures a very low complexity class, namely AC^0 , a class which is defined in terms of circuits; cf. [95].

The study of the complexity of evaluating a formula φ of a logic \mathcal{L} in a structure \mathcal{A} arises in various contexts. For example, \mathcal{A} may be a database instance and φ a corresponding query, or \mathcal{A} may represent the state space of a program and φ a desired property. When considering such evaluations the following kinds of complexities have been treated (the first one being the subject of this chapter):

- *data complexity* of \mathcal{L} : For a fixed sentence, we measure the complexity as a function of the size of the structure;
- *expression complexity* of \mathcal{L} : For a fixed structure, we measure the complexity as a function of the length of the formula;
- *combined complexity* of \mathcal{L} : It is measured as a function of both the size of the structure and the length of the formula.

In this book we mainly concentrate on data complexity. The following exercise contains two results on combined complexity. Part (a) illustrates a general phenomenon, namely a considerable increase against data complexity (cf. [52]); part (b) shows that it may pay to look for axiomatizations with few variables. For a proof the reader should carefully analyze the proof of 7.4.1.

Exercise 7.5.3 (a) The combined complexity of FO is in PSPACE. (In fact, it is PSPACE-complete; see [142].)

(b) For $s \geq 1$, the combined complexity of FO^s is in PTIME. (In fact, it is PTIME-complete for $s \geq 2$; see [143].) \square

The descriptive characterizations of complexity classes given by the main theorem are of importance in various respects:

- They may help to recognize that a concrete problem is in a given complexity class (by expressing it in the corresponding logic).
- They allow to view the logics involved as higher programming languages for problems of the corresponding complexity class. (Note that the proofs of the preceding section show how to convert a sentence φ into an algorithm accepting the class of models of φ and satisfying the required resource restrictions.)
- Characteristic features of the logic may be seen as characteristic features of the complexity class described by it and may add to a better understanding. (For instance, the result about $\text{FO}(\text{IFP})$ and PTIME shows us that inflationary inductions are an essential ingredient of PTIME algorithms.)
- The descriptive characterizations allow to convert problems, methods, and results of complexity theory into logic and vice versa, thus widening the methodological possibilities for both sides.

In the remainder of this section we give some illustrating examples for the last point. We first stay with ordered structures and subsequently discuss the role of order.

Sentences φ and ψ in a vocabulary τ with $< \in \tau$ are said to be *equivalent on ordered structures* if for all ordered τ -structures \mathcal{A} ,

$$\mathcal{A} \models \varphi \quad \text{iff} \quad \mathcal{A} \models \psi.$$

Corollary 7.5.4 *On ordered structures, every $\text{FO}(\text{IFP})$ -sentence is equivalent to an $\text{FO}(\text{IFP})$ -sentence in which IFP occurs at most once. The same applies to $\text{FO}(\text{PFP})$ and PFP.*

Proof. Suppose $\varphi \in \text{FO}(\text{IFP})[\tau]$ with $< \in \tau$. Then $\text{ordMod}(\varphi) \in \text{PTIME}$. Now the claim follows from 7.3.5. For $\text{FO}(\text{PFP})$ cf. 7.3.3. \square

The next corollary also has a simple direct proof.

Corollary 7.5.5 *Let \mathcal{C} be one of the complexity classes mentioned in the Main Theorem. If K is a class of ordered structures in \mathcal{C} then there is a Turing machine M strongly witnessing $K \in \mathcal{C}$, that is,*

- M accepts K
- every run stops in the accepting or in the rejecting state
- every run fulfills the time or space bounds characteristic for \mathcal{C} .

Proof. Let \mathcal{L} be the logic capturing \mathcal{C} . Then there is a sentence of \mathcal{L} axiomatizing K . By the results of section 7.4 we know that for every class K axiomatizable in \mathcal{L} there is a machine strongly witnessing $K \in \mathcal{C}$. \square

An immediate consequence of the Main Theorem is the equivalence of (i) and (ii):

- (i) $\text{PTIME} = \text{PSPACE}$
- (ii) $\text{FO}(\text{IFP}) \equiv \text{FO}(\text{PFP})$ on ordered structures.

Note, however, that here PTIME and PSPACE are understood as classes of ordered structures and not as languages over alphabets. Does (i) mean the same as $\text{PTIME} = \text{PSPACE}$ in complexity theory? We want to show this by making clear that here and in complexity theory we deal only with different presentations of a complexity class.

If \mathcal{C} is a complexity class of complexity theory, we denote by \mathcal{C}' the corresponding complexity class of structures. For example, for PTIME we have

\mathcal{C} consists of all languages L , $L \subseteq \mathbb{A}^+$ for some alphabet \mathbb{A} , such that there exists a deterministic Turing machine accepting L in polynomial time

and

\mathcal{C}' consists of all classes K , $K \subseteq \mathcal{O}[\tau]$ for some τ with $< \in \tau$, such that there is a deterministic Turing machine M accepting K in polynomial time.

In the following let $\mathcal{C}, \mathcal{C}_1, \mathcal{C}_2$ range over the complexity classes LOGSPACE , NLOGSPACE , PTIME , NPTIME , and PSPACE of complexity theory.

Our first consideration will show that $\mathcal{C} \subseteq \mathcal{C}'$ up to a natural transition from words to ordered structures, thereby using word models (already introduced in section 6.2).

Fix an alphabet \mathbb{A} and let $\tau(\mathbb{A})$ be the vocabulary $\{<\} \cup \{P_a \mid a \in \mathbb{A}\}$ with unary P_a . If $u \in \mathbb{A}^+$, denote by K_u the class of structures of the form

$$(B, <, (P_a)_{a \in \mathbb{A}}),$$

where the cardinality of B equals the length of u , $<$ is an ordering of B , and P_a corresponds to the positions in u carrying an a . For $L \subseteq \mathbb{A}^+$ set $K(L) := \bigcup_{u \in L} K_u$. Clearly, $K(L) \subseteq \mathcal{O}[\tau(\mathbb{A})]$ and, since

$$K(\mathbb{A}^+) = \text{ordMod}(\forall x (\bigvee_{a \in \mathbb{A}} (P_a x \wedge \bigwedge_{\substack{b \in \mathbb{A} \\ b \neq a}} \neg P_b x))),$$

we have

$$(1) \quad K(\mathbb{A}^+) \in \text{LOGSPACE}.$$

One can easily show that

$$(2) \quad \text{for } L \subseteq \mathbb{A}^+, \quad L \in \mathcal{C} \text{ iff } K(L) \in \mathcal{C}',$$

thus obtaining “ $\mathcal{C} \subseteq \mathcal{C}'$ up to transitions”.

Now we turn to a statement showing $\mathcal{C}' \subseteq \mathcal{C}$ up to a transition from ordered structures to words. Let τ with $< \in \tau$ be given. Set $\mathbb{A}_0 := \{0, 1, \alpha, \omega\}$. For $\mathcal{A} \in \mathcal{O}[\tau]$ let $u_{\mathcal{A}}$ be the word in \mathbb{A}_0^+ obtained by concatenating the inscriptions on all input tapes of a Turing machine started with input \mathcal{A} (where we include the “virtual letters” α and ω). For a class $K \subseteq \mathcal{O}[\tau]$ set

$$L(K) := \{u_{\mathcal{A}} \mid \mathcal{A} \in K\}.$$

Clearly, given τ , there is a polynomial $p \in \mathbb{N}[x]$ such that for all $\mathcal{A} \in \mathcal{O}[\tau]$ we have $\|\mathcal{A}\| \leq |u_{\mathcal{A}}| \leq p(\|\mathcal{A}\|)$. In particular, for $p(x) := x^d$ we have that $\log \|\mathcal{A}\| \leq \log |u_{\mathcal{A}}| \leq d \cdot \log \|\mathcal{A}\|$. Invoking these relations one shows

$$(3) \quad L(\mathcal{O}[\tau]) \in \text{LOGSPACE}$$

$$(4) \quad \text{for } K \subseteq \mathcal{O}[\tau], \quad K \in \mathcal{C}' \text{ iff } L(K) \in \mathcal{C}$$

((3) is immediate, and (4) is left as an exercise to the reader). Thus, “ $\mathcal{C}' \subseteq \mathcal{C}$ up to transitions”.

From (2) and (4) we infer

Proposition 7.5.6 $\mathcal{C}_1 \subseteq \mathcal{C}_2$ iff $\mathcal{C}'_1 \subseteq \mathcal{C}'_2$.

Proof. First, suppose $\mathcal{C}_1 \subseteq \mathcal{C}_2$ and let $K \in \mathcal{C}'_1$ where $K \subseteq \mathcal{O}[\tau]$. Then by (4), $L(K) \in \mathcal{C}_1$ and by hypothesis, $L(K) \in \mathcal{C}_2$. Therefore, $K \in \mathcal{C}'_2$ (again by (4)). Now assume $\mathcal{C}'_1 \subseteq \mathcal{C}'_2$ and let $L \in \mathcal{C}_1$. Then, $K(L) \in \mathcal{C}'_1$ by (2), and hence, $K(L) \in \mathcal{C}'_2$. Therefore, $L \in \mathcal{C}_2$ (again by (2)). \square

Together with the Main Theorem we have

Corollary 7.5.7 (a) $\text{FO(IFP)} \equiv \text{FO(PFP)}$ on ordered structures
iff $\text{PTIME} = \text{PSPACE}$ (in complexity theory).

- (b) $\text{FO}(\text{IFP}) \equiv \Sigma_1^1$ on ordered structures
 iff $\text{PTIME} = \text{NPTIME}$ (in complexity theory). \square

Moreover, the preceding argument shows that in order to get $\text{FO}(\text{IFP}) \equiv \text{FO}(\text{PFP})$ on ordered structures (or, equivalently, $\text{PTIME} = \text{PSPACE}$) it suffices to prove that $\text{FO}(\text{IFP}) \equiv \text{FO}(\text{PFP})$ holds for ordered τ -structures where, besides $<$, only unary relation symbols are in τ . Since one need only consider languages over $\{0, 1\}$, even a single unary relation symbol suffices.

Corollary 7.5.8 *The following are equivalent:*

- (i) $\text{PTIME} = \text{NPTIME}$
 (ii) $\text{FO}(\text{IFP}) \equiv \text{SO}$ on ordered structures.

Proof. If (ii) holds then $\Sigma_1^1 \leq \text{FO}(\text{IFP})$ on ordered structures, thus $\text{NPTIME} \leq \text{PTIME}$. Conversely, if $\text{NPTIME} = \text{PTIME}$ then, on ordered structures, $\Sigma_1^1 \equiv \text{FO}(\text{IFP})$. As Σ_1^1 is closed under existential quantifications and $\text{FO}(\text{IFP})$ under boolean operations, an easy induction yields $\text{SO} \equiv \text{FO}(\text{IFP})$. \square

Whereas the preceding corollaries contain the translation of *problems* from complexity theory to logics we now turn to the translation of a *result*. In complexity theory one shows

$$\text{LOGSPACE} \subseteq \text{NLOGSPACE} \subseteq \text{PTIME} \subseteq \text{NPTIME} \subseteq \text{PSPACE}$$

and

$$\text{LOGSPACE} \neq \text{PSPACE}.$$

Hence, by the Main Theorem,

Corollary 7.5.9 *On ordered structures,*

- (a) $\text{FO}(\text{DTC}) \leq \text{FO}(\text{TC}) \leq \text{FO}(\text{IFP}) \leq \Sigma_1^1 \leq \text{FO}(\text{PFP})$.
 (b) $\text{FO}(\text{DTC}) \not\equiv \text{FO}(\text{PFP})$. \square

Note that most of the \leq -relations in (a) are immediate; the remaining ones and part (b) will be obtained in chapters 8 and 9 by purely modeltheoretic means.

For any of the complexity classes \mathcal{C} introduced so far, in complexity theory one defines the class $\text{co-}\mathcal{C}$ to be the class of complements of languages in \mathcal{C} , that is, for any alphabet \mathbb{A} and $L \subseteq \mathbb{A}^+$,

$$L \in \text{co-}\mathcal{C} \quad \text{iff} \quad (\mathbb{A}^+ \setminus L) \in \mathcal{C}.$$

Clearly, any deterministic class \mathcal{C} is closed under complements, that is, $\mathcal{C} = \text{co-}\mathcal{C}$. Similarly, we define the class $\text{co-}\mathcal{C}'$ as the class of complements of classes of structures in \mathcal{C}' ; more precisely, for τ with $< \in \tau$ and $K \subseteq \mathcal{O}[\tau]$, we set

$$K \in \text{co-}\mathcal{C}' \quad \text{iff} \quad (\mathcal{O}[\tau] \setminus K) \in \mathcal{C}'.$$

Exercise 7.5.10 Show that Proposition 7.5.6 remains true if one allows \mathcal{C}_1 and \mathcal{C}_2 to range over the complexity classes mentioned in the Main Theorem and their complements. (Use (1)-(4) before 7.5.6.) \square

We now are going to discuss the role of order and to get information whether and to what extent orderings can be avoided. Let \mathcal{A} be a not necessarily ordered structure. We already mentioned that in order to consider \mathcal{A} as an input for a Turing machine, we have to represent it as a string (or a sequence of strings), for example, by labeling the elements of \mathcal{A} in some way. Taking, say, the lexicographic ordering of the labels, we get an ordering on A and hence, an ordered structure. Now, if \mathcal{A} is a graph we can state questions such as “Is there a path from the 5-th to the 28-th element?” whose answer depends on this ordering and is senseless for \mathcal{A} itself. The following framework enables us to concentrate on questions intrinsic to \mathcal{A} .

Definition 7.5.11 Let K be a class of (unordered) τ -structures. Set $\tau_{<} := \tau \dot{\cup} \{<\}$. The class $K_{<}$ of *ordered representations* of structures in K is given by

$$K_{<} := \{(\mathcal{A}, <) \mid \mathcal{A} \in K, < \text{ an ordering of } A\}. \quad \square$$

If \mathcal{L} is a logic capturing the complexity class \mathcal{C} we have

$$K_{<} \in \mathcal{C} \quad \text{iff} \quad \text{there is } \varphi \in \mathcal{L}[\tau_{<}] \text{ such that } K_{<} = \text{Mod}(\varphi).$$

The sentence φ on the right side of the equivalence is order-invariant in the finite, since for every \mathcal{A} and any orderings $<_1$ and $<_2$ of A we have

$$(\mathcal{A}, <_1) \in K_{<} \quad \text{iff} \quad (\mathcal{A}, <_2) \in K_{<}$$

and thus,

$$(\mathcal{A}, <_1) \models \varphi \quad \text{iff} \quad (\mathcal{A}, <_2) \models \varphi.$$

Now, if \mathcal{L} would be closed under order-invariant sentences in the finite (by 3.5.2, FO does not have this property), we would have

$$K_{<} \in \mathcal{C} \quad \text{iff} \quad \text{there is } \psi \in \mathcal{L}[\tau] \text{ such that } K = \text{Mod}^\tau(\psi).^8$$

In general, this does not hold. To give a counterexample for FO(DTC), let $K = \text{EVEN}[\tau]$ with $\tau = \emptyset$ be the class of sets of even cardinality. Since $K_{<} \in \text{LOGSPACE}$, there is a sentence φ of FO(DTC)[$\tau_{<}$] such that

$$K_{<} = \text{Mod}(\varphi);$$

for example, as φ we can take the sentence $\neg[\text{DTC}_{x,y} y = x + 2] \min \max$, where we use self-explanatory abbreviations. This sentence is order-invariant

⁸ We use $\text{Mod}^\tau(\dots)$ with upper index τ to rule out any ambiguity on the vocabulary considered.

in the finite: the evaluation of φ in a structure $(A, <^A)$ makes use of the ordering $<^A$, but the outcome of this evaluation does not depend on the specific ordering $<^A$ we have chosen. We shall see in 8.4.4 that for no sentence ψ of $\text{FO}(\text{DTC})[\tau]$, even of $\text{FO}(\text{PFP})[\tau]$,

$$K = \text{Mod}(\psi).$$

These observations lead to a stronger notion of what it means to capture a complexity class.

Definition 7.5.12 Let \mathcal{L} be a logic and \mathcal{C} a complexity class. \mathcal{L} *strongly captures* \mathcal{C} if for all vocabularies τ and all classes K of τ -structures,

$$K_{<} \in \mathcal{C} \text{ iff } K \text{ is axiomatizable in } \mathcal{L}. \quad \square$$

The following proposition holds for all complexity classes \mathcal{C} considered so far; essentially one needs that \mathcal{C} contains LOGSPACE. The proof is left to the reader as an exercise.

Proposition 7.5.13 *If \mathcal{L} strongly captures \mathcal{C} then \mathcal{L} captures \mathcal{C} .* □

The converse is false: The counterexample given before the definition shows that $\text{FO}(\text{PFP})$ does not capture PSPACE strongly. For the class $\text{EVEN}[\tau]$ we used as a counterexample, we have $\text{EVEN}[\tau]_{<} \in \text{LOGSPACE}$, and $\text{EVEN}[\tau]$ is not axiomatizable in $\text{FO}(\text{PFP})$. Since $\text{FO}(\text{DTC}) \leq \text{FO}(\text{TC}) \leq \text{FO}(\text{IFP}) \leq \text{FO}(\text{PFP})$ holds for arbitrary structures, we see that none of these logics strongly captures the complexity class corresponding to it by the Main Theorem. The result cannot be extended to Σ_1^1 and NPTIME, since we know $\Sigma_1^1 \leq \text{FO}(\text{PFP})$ only on *ordered* structures. In fact, we have

Theorem 7.5.14 Σ_1^1 *strongly captures* NPTIME.

Proof. Let τ be arbitrary and K be a class of τ -structures. Assume $K = \text{Mod}^\tau(\varphi)$ for some $\Sigma_1^1[\tau]$ -sentence φ , say, $\varphi = \exists X_1 \dots \exists X_m \psi$ with first-order ψ . Set $\chi := \exists X_1 \dots \exists X_m (\psi \wedge \text{“} < \text{ is an ordering”})$. Then, $\chi \in \Sigma_1^1[\tau_{<}]$, $\text{Mod}(\chi) = K_{<}$, and hence, $K_{<} \in \text{NPTIME}$ by the Main Theorem. Conversely, if $K_{<} \in \text{NPTIME}$ then, again by the Main Theorem, there is a sentence $\varphi \in \Sigma_1^1[\tau_{<}]$ such that $K_{<} = \text{Mod}^{\tau_{<}}(\varphi)$. Set $\psi := \exists < \varphi$. Then $\psi \in \Sigma_1^1[\tau]$, and for any τ -structure \mathcal{A} we have

$$\begin{aligned} \mathcal{A} \models \psi & \text{ iff } \text{there is } <^A \text{ with } (\mathcal{A}, <^A) \models \varphi \\ & \text{ iff } \text{there is } <^A \text{ with } (\mathcal{A}, <^A) \in K_{<} \\ & \text{ iff } \mathcal{A} \in K, \end{aligned}$$

that is, $K = \text{Mod}^\tau(\psi)$. □

Theorem 7.5.15 Π_1^1 *strongly captures* co-NPTIME.

Proof. Let τ be arbitrary and K be a class of τ -structures. If $\text{Str}[\tau]$ denotes the class of all τ -structures then

$$(+) \quad (\text{Str}[\tau] \setminus K)_{<} = \mathcal{O}[\tau_{<}] \setminus K_{<}$$

Therefore,

$$\begin{aligned} K \in \Pi_1^1 & \text{ iff } \text{there is } \psi(X_1, \dots, X_m) \in \text{FO}[\tau] \text{ such that} \\ & \quad K = \text{Mod}(\forall X_1 \dots \forall X_m \psi) \\ & \text{ iff } \text{there is } \chi(X_1, \dots, X_m) \in \text{FO}[\tau] \text{ such that} \\ & \quad \text{Str}[\tau] \setminus K = \text{Mod}(\exists X_1 \dots \exists X_m \chi) \\ & \text{ iff } (\text{Str}[\tau] \setminus K)_{<} \in \text{NPTIME} \quad (\text{by 7.5.14}) \\ & \text{ iff } (\mathcal{O}[\tau_{<}] \setminus K_{<}) \in \text{NPTIME} \quad (\text{by } (+)) \\ & \text{ iff } K_{<} \in \text{co-NPTIME}. \end{aligned}$$

□

Corollary 7.5.16 $\text{NPTIME} = \text{co-NPTIME}$ iff $\Sigma_1^1 \equiv \Pi_1^1$.

Proof. As an example we show the implication from left to right. Let K be a class of structures. Then

$$\begin{aligned} K \in \Sigma_1^1 & \text{ iff } K_{<} \in \text{NPTIME} \\ & \text{ iff } K_{<} \in \text{co-NPTIME} \quad (\text{by hypothesis}) \\ & \text{ iff } K \in \Pi_1^1. \end{aligned}$$

□

Corollary 7.5.17 $\text{NPTIME} = \text{co-NPTIME}$ iff $\text{SO} \equiv \Sigma_1^1$.

Proof. By the preceding corollary it suffices to prove

$$\Pi_1^1 \equiv \Sigma_1^1 \quad \text{iff} \quad \text{SO} \equiv \Sigma_1^1.$$

For a logic \mathcal{L} we write $\varphi \underset{\sim}{\in} \mathcal{L}$ to express that the sentence φ is equivalent to an \mathcal{L} -sentence. Clearly,

$$\begin{aligned} \varphi \in \Sigma_1^1 & \text{ implies } \neg\varphi \underset{\sim}{\in} \Pi_1^1, \\ \varphi \in \Pi_1^1 & \text{ implies } \neg\varphi \underset{\sim}{\in} \Sigma_1^1. \end{aligned}$$

Now suppose that $\text{SO} \equiv \Sigma_1^1$ and therefore, $\Pi_1^1 \leq \Sigma_1^1$. Let $\varphi \in \Sigma_1^1$. Then, $\neg\varphi \underset{\sim}{\in} \Pi_1^1$ and hence, $\neg\varphi \underset{\sim}{\in} \Sigma_1^1$. Therefore, $\varphi \underset{\sim}{\in} \Pi_1^1$. Conversely, assume that $\Pi_1^1 \equiv \Sigma_1^1$. One easily shows that the class Σ_1^1 is closed – up to equivalence – under \forall and existential first-order and second-order quantifications. For closure under \neg argue as follows: Suppose $\varphi = \neg\psi$ and $\psi \in \Sigma_1^1$. Then $\neg\psi \underset{\sim}{\in} \Pi_1^1 \equiv \Sigma_1^1$.

□

Exercise 7.5.18 Show for $k \geq 1$ that $\Sigma_k^1 \equiv \Sigma_{k+1}^1$ implies $\Sigma_k^1 \equiv \text{SO}$. The question, whether $\Sigma_k^1 \equiv \Sigma_{k+1}^1$ for some k , is open. However, it is conjectured that

$$\Sigma_1^1 < \Sigma_2^1 < \dots$$

For $k \geq 1$, the class of ordered structures axiomatizable in Σ_k^1 forms the k -th stage PH_k of the *polynomial hierarchy* PH. Its initial stage PH_0 is defined to be PTIME. In 12.4.8 we shall give a definition of PH in terms of complexity. \square

Exercise 7.5.19 Define the logic $\text{SO}(\text{PFP})$. Show that it captures PSPACE strongly. In 8.5.1 we shall implicitly show by a purely modeltheoretic proof that $\text{SO}(\text{PFP}) \equiv \text{FO}(\text{PFP})$ on ordered structures. \square

The question whether LOGSPACE, NLOGSPACE, and PTIME can be strongly captured by a logic is a major issue of descriptive complexity theory. Of course, one has to give a precise definition of what is meant by a “logic” in this context. Moreover, to rule out trivial positive answers, one has to require some conditions of effectiveness. Note that in the capturing results obtained in the preceding sections, we effectively assigned to every sentence a Turing machine accepting its models. And to every Turing machine accepting a class of models we effectively assigned a sentence axiomatizing this class. Up to now all attempts to find a logic that strongly captures PTIME in such an effective way have failed. In Chapter 11 we shall come back to such attempts.

We close this section by a result which completes the proof of the Main Theorem (compare 8.6.15 for a modeltheoretic proof).

Theorem 7.5.20 *On ordered structures, $\text{FO}(\text{posTC}) \equiv \text{FO}(\text{TC})$.*

Proof. We make use of the fact (shown in 8.6.14) that, on ordered structures, every formula of $\text{FO}(\text{posTC})$ is equivalent to a formula of the form

$$[\text{TC}_{\bar{x}, \bar{y}} \psi] \widetilde{\text{minmax}}$$

with first-order ψ .

The proof of the theorem proceeds by induction on $\text{FO}(\text{TC})$ -formulas, the only nontrivial case being the negation step. By the induction hypothesis and the fact just mentioned we may assume that

$$\varphi = \neg [\text{TC}_{\bar{x}, \bar{y}} \psi] \bar{s}\bar{t}$$

with first-order ψ . For simplicity, we assume that $\psi = \psi(\bar{x}, \bar{y})$. By the Main Theorem (more precisely, by 7.4.1) there is a Turing machine M_0 strongly witnessing that $\text{ordMod}(\psi) \in \text{LOGSPACE}$.

Suppose $\bar{x} = x_1 \dots x_r$. Given a structure \mathcal{A} and $\bar{a}, \bar{b} \in A^r$, let $d_\psi(\bar{a}, \bar{b})$ be the length of the shortest ψ -path connecting \bar{a} and \bar{b} ,

$$d_\psi(\bar{a}, \bar{b}) := \min\{k > 0 \mid \text{there exist } \bar{a}_0 = \bar{a}, \bar{a}_1, \dots, \bar{a}_k = \bar{b} \text{ such that} \\ \mathcal{A} \models \psi[\bar{a}_i, \bar{a}_{i+1}] \text{ for } i < k\},$$

where $d_\psi(\bar{a}, \bar{b}) := \infty$ in case the set on the right side is empty. Note that

- if $d_\psi(\bar{a}, \bar{b}) < \infty$ then $0 < d_\psi(\bar{a}, \bar{b}) \leq \|A\|^r$
- $d_\psi(\bar{a}, \bar{c}) \leq d_\psi(\bar{a}, \bar{b}) + d_\psi(\bar{b}, \bar{c})$.

Moreover, $\neg[\text{TC}_{\bar{x}, \bar{y}} \psi(\bar{x}, \bar{y})] \bar{s} \bar{t}$ is equivalent to

$$“\|\{\bar{v} \mid d_\psi(\bar{s}, \bar{v}) < \infty\}\| = \|\{\bar{v} \mid d_{(\psi(\bar{x}, \bar{y}) \wedge \neg \bar{y} = \bar{t})}(\bar{s}, \bar{v}) < \infty\}\|”.$$

We first show that there is a nondeterministic log space-bounded⁹ machine M such that for any ordered structure $(\mathcal{A}, \bar{a}, \bar{l}, \bar{w}, \bar{w}')$,

if $\|\{\bar{e} \mid d_\psi(\bar{a}, \bar{e}) \leq \bar{l}\}\| = \bar{w}$ then

$$M \text{ accepts } (\mathcal{A}, \bar{a}, \bar{l}, \bar{w}, \bar{w}') \text{ iff } \|\{\bar{e} \mid d_\psi(\bar{a}, \bar{e}) \leq \bar{l} + 1\}\| = \bar{w}',$$

where the corresponding natural numbers $\leq \|A\|^r$ are given by their $\|A\|$ -adic representations $\bar{l} = l_0 \dots l_r$, $\bar{w} = w_0 \dots w_r$, $\bar{w}' = w'_0 \dots w'_r$.

We present the machine M . When during its computation M checks whether $\mathcal{A} \models \psi[\bar{c}, \bar{d}]$ holds or not, this is done by invoking M_0 . Started with $(\mathcal{A}, \bar{a}, \bar{l}, \bar{w}, \bar{w}')$, M first sets a counter to \bar{w}' . Then, for every $\bar{b} \in A^r$, it carries out either (1) or (2), the choice being done nondeterministically:

- (1) M nondeterministically guesses a path witnessing $d_\psi(\bar{a}, \bar{b}) \leq \bar{l} + 1$ and decreases the counter by one; in case the counter is zero it rejects.
- (2) Using an additional counter with the initial value \bar{w} , M nondeterministically guesses \bar{w} many distinct tuples $\bar{c} \in A^r$ together with a proof that $d_\psi(\bar{a}, \bar{c}) \leq \bar{l}$; for each such \bar{c} it shows that $\bar{c} \neq \bar{b}$ and $\neg\psi[\bar{c}, \bar{b}]$ (thus, in case $\|\{\bar{e} \mid d_\psi(\bar{a}, \bar{e}) \leq \bar{l}\}\| = \bar{w}$, proving that $d_\psi(\bar{a}, \bar{b}) > \bar{l} + 1$). In case $\bar{l} = 0$ it shows $\neg\psi[\bar{a}, \bar{b}]$.

Finally, if all $\bar{b} \in A^r$ have been dealt with and the counter is 0, M accepts.

Since M is log space-bounded there is, by the Main Theorem, a formula $\chi_\psi(\bar{u}, \bar{y}, \bar{w}, \bar{w}') \in \text{FO}(\text{posTC})$ axiomatizing the class accepted by M . Then, noting that $\underbrace{1 \underline{0} \dots 0}_r$ is (the representation of) $\|A\|^r$, we get a formula

$\rho_\psi(\bar{u}, \bar{w}) \in \text{FO}(\text{posTC})$ with the meaning

$$“\|\{\bar{v} \mid d_\psi(\bar{u}, \bar{v}) < \infty\}\| = \bar{w}”$$

or, equivalently, with the meaning

$$“\|\{\bar{v} \mid d_\psi(\bar{u}, \bar{v}) \leq \|A\|^r\}\| = \bar{w}”$$

⁹ log space-bounded means $c \cdot \log$ space-bounded for some $c > 0$.

by setting

$$\rho_\psi(\bar{u}, \bar{w}) := [\text{TC}_{\bar{y}\bar{w}, \bar{y}'\bar{w}'}(\chi_\psi(\bar{u}, \bar{y}, \bar{w}, \bar{w}') \wedge \bar{y}' = \bar{y} + 1)] \underbrace{\text{minmin} 1 0 \dots 0}_{r \text{ times}} \bar{w}.$$

Since we saw above that $\varphi = \neg [\text{TC}_{\bar{x}, \bar{y}} \psi(\bar{x}, \bar{y})] \bar{s} \bar{t}$ is equivalent to

$$\|\{\bar{v} \mid d_\psi(\bar{s}, \bar{v}) < \infty\}\| = \|\{\bar{v} \mid d_{(\psi(\bar{x}, \bar{y}) \wedge \neg \bar{y} = \bar{t})}(\bar{s}, \bar{v}) < \infty\}\|,$$

we obtain that φ is equivalent to the formula

$$\exists \bar{z} (\rho_\psi(\bar{s}, \bar{z}) \wedge \rho_{(\psi \wedge \neg \bar{y} = \bar{t})}(\bar{s}, \bar{z})),$$

a formula of FO(posTC). □

As a consequence of part (b) of the Main Theorem whose proof we just have completed we get, since FO(TC) is closed under negation:

Corollary 7.5.21 NLOGSPACE = co-NLOGSPACE. □

In the preceding considerations, when treating questions of effectiveness, we have restricted ourselves to ordered structures \mathcal{A} , as they offer a natural way of encoding by passing to the unique isomorphic copy whose universe is $\{0, \dots, \|\mathcal{A}\| - 1\}$ and whose order is the natural order on $\{0, \dots, \|\mathcal{A}\| - 1\}$. There is an alternative way, namely the direct restriction to structures whose universe forms an initial segment of the natural numbers. Of course, both frameworks are equivalent:

Exercise 7.5.22 Let τ be an arbitrary vocabulary. A *numerical* τ -structure is a τ -structure whose universe is an initial segment of the natural numbers, that is, the universe is the set $\{0, \dots, n\}$ for some n . For a class K of τ -structures let

$$K_{\text{num}} := \{\mathcal{A} \in K \mid \mathcal{A} \text{ is numerical}\}$$

be the set of numerical structures in K , and for a sentence φ let

$$\text{Mod}_{\text{num}}(\varphi) := \{\mathcal{A} \mid \mathcal{A} \models \varphi, \mathcal{A} \text{ numerical}\}$$

be the set of its numerical models.

Let \mathcal{C} and \mathcal{L} be any of the complexity classes and logics, respectively, considered so far. Show

- (a) For any class K , we have $K_{<} \in \mathcal{C}$ iff $K_{\text{num}} \in \mathcal{C}$ (that is, there is a Turing machine according to \mathcal{C} that, started with any numerical \mathcal{A} accepts \mathcal{A} iff $\mathcal{A} \in K_{\text{num}}$).
- (b) \mathcal{L} strongly captures \mathcal{C} iff for any τ and any class K of τ -structures, $K_{\text{num}} \in \mathcal{C}$ iff there is a sentence φ of $\mathcal{L}[\tau]$ such that $K_{\text{num}} = \text{Mod}_{\text{num}}(\varphi)$.
- (c) \mathcal{L} captures \mathcal{C} iff for any τ and any class K of τ -structures, $K_{\text{num}} \in \mathcal{C}$ iff there is an \mathcal{L} -sentence φ of $\mathcal{L}[\tau_{<}]$ with $K_{\text{num}} = \{\mathcal{A} \mid (\mathcal{A}, <^{\mathcal{A}}) \models \varphi, \mathcal{A} = \{0, \dots, n\} \text{ for some } n, \text{ and } <^{\mathcal{A}} \text{ is the natural ordering on } \mathcal{A}\}$. □

7.5.1 Appendix¹⁰

In the preceding sections Turing machines have only been used as *acceptors*, that is, given a machine M and an input structure \mathcal{A} , the question was whether M would accept \mathcal{A} . However, Turing machines can also be used to calculate functions. For later purposes we are interested in machines that map structures to structures, so-called transducers.

Let τ and σ be vocabularies containing $<$. We consider (τ, σ) -Turing machines, that is, deterministic Turing machines M that have input tapes for encoding τ -structures, work tapes, and in addition so-called *output tapes*, namely as many as we need for our way of encoding σ -structures as inputs. The output tapes are write-only tapes: In the beginning of a calculation they are empty, and their heads scan the first cell. In a step of M the heads on the output tapes either are not active or write a letter on the scanned cell and then move to the right neighbour. (So the final positions of the heads on the output tapes take over the role of the endmark ω since they determine the portion of the tape visited by the heads during the computation.) Moreover, the behaviour of M does not depend on the actual inscriptions and head positions of the output tapes. Altogether it is made precise by the form of the instructions that look like (cf. $(*)$ in subsection 7.2.2):

$$(+) \quad sb_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s' c'_1 \dots c'_m h_0 \dots h_{k+l+m} d_0 \dots d_p$$

(in case M has $p + 1$ output tapes OP_0, \dots, OP_p), where $d_i \in \{00, 10, 11\}$. The meaning of $(+)$ is fixed as for $(*)$ in subsection 7.2.2 adding that

$d_i = 00$	means	“Do nothing on OP_i ”,
$d_i = 10$	means	“Write 0 on the scanned cell of OP_i and move the head to the cell on the right”,
$d_i = 11$	means	“Write 1 on the scanned cell of OP_i and move the head to the cell on the right”.

A (τ, σ) -Turing machine M is a (τ, σ) -transducer if M , started with an ordered τ -structure \mathcal{A} , finally stops and the output tapes then carry the encoding of an ordered σ -structure \mathcal{B} (where, up to the missing endmark ω , the same encoding as for input structures is used). To be definite, we let $\{0, 1, \dots, \|B\| - 1\}$ be the universe of \mathcal{B} .

When considering space-bounded transducers, the space used on the output tapes (similar to that used on the input tapes) does not count. Hence, a log space-bounded transducer may fill a number of cells on the output tapes that is not log space-bounded in the size of the input.

Without changing the space needed and, up to a constant factor, the time needed, we may assume that transducers are *normalized* in the sense that in the instructions $(+)$ either (i) or (ii), where

¹⁰ The notions and results of this appendix are only needed in Chapters 11 and 12. They may be skipped in a first reading.

- (i) all d_i are = 00
- (ii) exactly one d_i is $\neq 00$, $c_1 \dots c_m = c'_1 \dots c'_m$, and $h_0 = \dots = h_{k+l+m} = 0$.

So in a step of a normalized transducer either no output tape is involved or only one output tape, but no other tape. In case (ii), we speak of an $(i, 0)$ -instruction if $d_i = 10$, and of an $(i, 1)$ -instruction if $d_i = 11$.

Fix τ, σ as above and let M be a normalized (τ, σ) -transducer. If we forget the output tapes and the d_i , we get a usual deterministic Turing machine M' . Thus, if

$$(1) \quad sb_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s'c'_1 \dots c'_m h_0 \dots h_{k+l+m} d_0 \dots d_p$$

is an instruction of M then

$$(2) \quad s_0 b_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s'c_1 \dots c'_m h_0 \dots h_{k+l+m}$$

is an instruction of M' ; we call it an (i, j) -instruction of M' if (1) is an (i, j) -instruction of M . When carrying out an (i, j) -instruction, M' at most changes its state.

Now, start M' with an ordered τ -structure \mathcal{A} . Then, for every point of the run of M started with \mathcal{A} , the inscription of the output tapes of M (and hence, the output structure of M) can be reconstructed from the run of M' as follows:

OP_i has 0 in cell l if the corresponding sequence of configurations of M' contains at least l many $(i, 0)$ - or $(i, 1)$ -configurations, the l -th one being an $(i, 0)$ -configuration.¹¹ And similarly, OP_i has 1 in cell l if the l -th $(i, 0)$ - or $(i, 1)$ -configuration is an $(i, 1)$ -configuration.

This observation enables us to describe the behaviour of M and, in particular, the output structure via the behaviour of M' , using formulas of a suitable logic as we have done it for acceptors in section 7.3. As an example that we shall need later we consider the case that M is log space-bounded and the logic is FO(DTC), recalling that the behaviour of log space-bounded acceptors can be described by formulas of FO(DTC) (cf. 7.3.8).

Theorem 7.5.23 *Let M be a log space-bounded (τ, σ) -transducer. There are an $e \geq 1$, an FO(DTC)-formula $\varphi_{\text{uni}}(x_1, \dots, x_e)$, and for every, say r -ary, $R \in \sigma$ an FO(DTC)-formula $\varphi_R(\bar{x}_1, \dots, \bar{x}_r)$ where $\text{length}(\bar{x}_1) = \dots = \text{length}(\bar{x}_r) = e$ such that for every ordered τ -structure \mathcal{A} the machine M , started with \mathcal{A} , leads to an output structure isomorphic to*

$$(\chi_{\text{uni}}^{\mathcal{A}}(-), (\varphi_R^{\mathcal{A}}(-, \dots, -))_{R \in \sigma}).$$

We leave the details of a proof to the reader. One uses an FO(DTC)-description of the run of M . The following exercise contains the key parts

¹¹ An (i, j) -configuration is a configuration calling an (i, j) -instruction.

needed to obtain the formulas mentioned in the theorem. Note that, when started with an ordered structure \mathcal{A} , both M and M' stop after the same number of steps. As the number of different, say, $d \cdot \log \|A\|$ -bounded configurations of M' is polynomial in $\|A\|$, both M and M' are $(n^e - 1)$ time-bounded for a suitable e . We may assume that $e = d$.

For notational simplicity let $\sigma = \{E, <\}$ with binary E .

Exercise 7.5.24 Referring to M' and using the formulas $\chi_{\text{start}}(\bar{x})$ and $\chi_{\text{succ}}(\bar{x}, \bar{x}')$ from Lemma 7.3.7, show for sufficiently large τ -structures \mathcal{A} :

(a) For $i \leq p$ there is an FO(DTC)-formula $\chi_{i\text{-free}}(\bar{x}, \bar{x}')$ such that for \bar{a}, \bar{a}' in A :

$$\mathcal{A} \models \chi_{i\text{-free}}[\bar{a}, \bar{a}'] \quad \text{iff} \quad \bar{a} \text{ and } \bar{a}' \text{ are } (d \cdot \log \|A\|)\text{-bounded configurations of } M', \text{ and } M' \text{, with } \mathcal{A} \text{ written on its input tapes, runs from } \bar{a} \text{ to } \bar{a}' \text{ without passing an } (i, 0)\text{-configuration or an } (i, 1)\text{-configuration.}$$

(b) For $i \leq p$ there is an FO(DTC)-formula $\chi_{i\text{-count}}(\bar{x}, \bar{y})$ with $\text{length}(\bar{y}) = e$ such that for \bar{a}, \bar{b} in A :

$$\mathcal{A} \models \chi_{i\text{-count}}[\bar{a}, \bar{b}] \quad \text{iff} \quad \bar{a} \text{ is the } \bar{b}\text{-th } (i, 0)\text{-configuration or } (i, 1)\text{-configuration of } M' \text{ started with } A \text{ (we refer to the lexicographical ordering of } A^e\text{).}$$

(c) Let $\mathcal{B} = (B, E^B, <^B)$ be the output structure that M gives on \mathcal{A} . There are FO(DTC)-formulas $\chi_{\text{uni}}(\bar{y}), \chi_E(\bar{y}, \bar{z}),$ and $\chi_{<}(\bar{y}, \bar{z})$ with $\text{length}(\bar{y}) = \text{length}(\bar{z}) = e$ such that

$$\begin{aligned} B &= \{l \mid \mathcal{A} \models \chi_{\text{uni}}[\bar{b}], \bar{b} \text{ is the } l\text{-th element in the lexicographical ordering of } A^e\}; \\ E^B &= \{(l, m) \mid \mathcal{A} \models \chi_E[\bar{b}, \bar{c}], \bar{b} \text{ is the } l\text{-th and } \bar{c} \text{ the } m\text{-th element in the lexicographical ordering of } A^e\}; \\ <^B &= \{(l, m) \mid \mathcal{A} \models \chi_{<}[\bar{b}, \bar{c}], \bar{b} \text{ is the } l\text{-th, } \bar{c} \text{ the } m\text{-th element}\} \text{ is the natural ordering on } B. \end{aligned}$$

□

Exercise 7.5.25 Formulate and prove the analogue of Theorem 7.5.23 for polynomially time-bounded (τ, σ) -transducers and FO(IFP). □

Notes 7.5.26 As a reference for computation and complexity theory we mention the books [87, 128, 46]. Theorem 6.2.1 goes back to [139]. The characterizations of complexity classes given by 7.5.2 are due to Immerman [93, 94] (LOGSPACE, NLOGSPACE), Immerman [91], Livchak [119], and Vardi [142] (PTIME), Fagin [36] (NPTIME), Abiteboul and Vianu [2] (PSPACE). Corollary 7.5.21 was independently proved in [94] and [137]. An important topic we have not treated here is circuit complexity. We refer the reader to [95] for information and further references.

8. Logics with Fixed-Point Operators

In Chapter 7 we have introduced logics with fixed-point operators in the context of computational complexity and shown that they capture important complexity classes. In the present chapter we study the (finite) model theory of fixed-point logics. Though we sometimes refer to Chapter 7, we repeat the relevant definitions and results to provide an independent approach.

In the rest of the book all structures will be finite, unless stated otherwise explicitly. Equivalence of formulas means equivalence with respect to all finite structures.

8.1 Inflationary and Least Fixed-Points

We prove some basic facts about fixed-points and, besides inflationary fixed-point logic, present an important sublogic, namely least fixed-point logic. We state two major results, the proofs being postponed to the next section: Inflationary and least fixed-point logic have the same expressive power, and every formula is equivalent to a formula containing at most one fixed-point application – a fact that extends Corollary 7.5.4 from ordered structures to arbitrary ones.

We observed in Chapter 2 that many important notions and global relations cannot be expressed by first-order formulas. An example is given by the reflexive and transitive closure of the edge relation in a graph (cf. 2.3.8). In the vocabulary $\tau := \{E\}$ for graphs consider the formula

$$\chi(x, y, X) := (x = y \vee \exists z(Xxz \wedge Ezy)).$$

It gives rise to a sequence of sets defined by

$$(*) \quad X_0 := \emptyset, \quad X_{n+1} := \{(x, y) \mid \chi(x, y, X_n)\}.$$

In other terms: Let $\mathcal{G} = (G, E^G)$ be a graph. Look at the function F^χ on the power set of $G \times G$, $F^\chi : \text{Pow}(G^2) \rightarrow \text{Pow}(G^2)$, given by

$$F^\chi(U) := \{(a, b) \in G^2 \mid \mathcal{G} \models \chi[a, b, U]\}$$

for $U \subseteq G \times G$. Then, $(*)$ corresponds to the sequence of sets

$$F_0^X := \emptyset, \quad F_{n+1}^X := F^X(F_n^X).$$

Note that

$$F_n^X = \{(a, b) \in G^2 \mid d(a, b) < n\},$$

where $d(a, b)$ denotes the distance between a and b in \mathcal{G} (see 1.A1). For $F_\infty^X := \bigcup_{n \geq 0} F_n^X$ we have $F(F_\infty^X) = F_\infty^X$, that is, F_∞^X is a fixed-point of F^X and

$$F_\infty^X = \{(a, b) \in G^2 \mid d(a, b) < \infty\}.$$

In particular, for $a \neq b$,

$$a, b \text{ are connected by a path in } \mathcal{G} \quad \text{iff} \quad (a, b) \in F_\infty^X,$$

and

$$\mathcal{G} \text{ is connected} \quad \text{iff} \quad F_\infty^X = G \times G.$$

The relations F_n^X that lead to the fixed-point F_∞^X are first-order definable, but in general, F_∞^X is not. In the fixed-point logics we are going to introduce the relation F_∞^X will be definable, too; in particular, the connectivity of graphs will be expressible.

Let us first study some aspects of the process above on a more abstract level. Fix a finite set M . A function $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ gives rise to a sequence of sets

$$\emptyset, F(\emptyset), F(F(\emptyset)), \dots$$

Denote its members by F_0, F_1, \dots , i.e., $F_0 = \emptyset$ and $F_{n+1} = F(F_n)$. F_n is called the n -th *stage* of F . Suppose that there is an $n_0 \in \mathbb{N}$ such that $F_{n_0+1} = F_{n_0}$, that is, $F(F_{n_0}) = F_{n_0}$. Then, $F_m = F_{n_0}$ for all $m \geq n_0$. We denote F_{n_0} by F_∞ and say that the *fixed-point* F_∞ of F *exists*. In case the fixed-point F_∞ does not exist, we agree to set $F_\infty := \emptyset$.

F is *inductive* if $F_0 \subseteq F_1 \subseteq \dots$

Lemma 8.1.1 (a) *If F_∞ exists then $F_\infty = F_{2\|M\|-1}$.*

(b) *If $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ is inductive then F_∞ exists and $F_\infty = F_{\|M\|}$.*

Proof. Part (a) coincides with (b) of 7.1.1.

(b) By assumption, $F_0 \subseteq F_1 \subseteq \dots \subseteq M$. Since M has $\|M\|$ elements, this sequence must get constant not later than with $F_{\|M\|}$. \square

F is *inflationary* if for all $X \subseteq M$

$$X \subseteq F(X),$$

and *monotone* if for all $X, Y \subseteq M$

$$X \subseteq Y \text{ implies } F(X) \subseteq F(Y).$$

- Lemma 8.1.2** (a) If F is inflationary or monotone then F is inductive.
 (b) If F is monotone then F_∞ is the least fixed-point of F , i.e., $F(F_\infty) = F_\infty$ and $F(X) = X$ implies $F_\infty \subseteq X$; even $F(X) \subseteq X$ implies $F_\infty \subseteq X$.
 (c) If $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ is arbitrary and $F' : \text{Pow}(M) \rightarrow \text{Pow}(M)$ is given by $F'(X) := X \cup F(X)$, then F' is inflationary. In case F is inductive we have $F'_n = F_n$ for all $n \geq 0$ and hence, $F'_\infty = F_\infty$.

Proof. (a) If F is inflationary then $F_n \subseteq F(F_n) = F_{n+1}$; thus F is inductive. Suppose that F is monotone. We have $F_0 = \emptyset \subseteq F(\emptyset) = F_1$ and therefore, by monotonicity, $F_1 = F(F_0) \subseteq F(F_1) = F_2$. Going on this way we obtain $F_n \subseteq F_{n+1}$ for all n .

(b) By part (a) and 8.1.1(b) the fixed-point F_∞ exists. Suppose $F(X) \subseteq X$. We have $F_0 = \emptyset \subseteq X$. Suppose $F_n \subseteq X$. Then, by monotonicity, $F_{n+1} = F(F_n) \subseteq F(X) \subseteq X$. Therefore, $F_\infty \subseteq X$.

(c) Clearly, F' is inflationary. Let F be inductive. We show by induction on n that $F'_n = F_n$. Obviously, $F'_0 = F_0$. Suppose $F'_n = F_n$. Then $F'_{n+1} = F'_n \cup F(F'_n) = F_n \cup F(F_n) = F_n \cup F_{n+1} = F_{n+1}$. \square

Exercise 8.1.3 Let $M := \{0, 1, 2\}$ and $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$. Show:

- (a) If $F(X) = \emptyset$ for all $X \subseteq M$, then F is monotone but not inflationary.
 (b) If

$$F(X) = \begin{cases} X & \text{for } X = M \\ X \cup \{\|X\|\} & \text{otherwise} \end{cases}$$

then F is inflationary but not monotone.

- (c) If

$$F(X) = \begin{cases} X & \text{for } X = \emptyset \\ X \setminus \{\|X\|\} & \text{otherwise} \end{cases}$$

then F is inductive but neither inflationary nor monotone. \square

Exercise 8.1.4 Let $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ be monotone. Show that F has a *greatest fixed-point*, i.e., there is $F_g \subseteq M$ such that

- $F(F_g) = F_g$
- $F(Y) = Y$ implies $Y \subseteq F_g$.

Define $G : \text{Pow}(M) \rightarrow \text{Pow}(M)$ by $G(X) := M \setminus F(M \setminus X)$. Prove that

- G is monotone
- $F_g = M \setminus G_\infty$, i.e., the greatest fixed-point of F is the complement of the least fixed-point of G . \square

Exercise 8.1.5 (a) Assume $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$ is *antitone*, that is, $X \subseteq Y$ implies $F(Y) \subseteq F(X)$. Show that

$$F_0 \subseteq F_2 \subseteq F_4 \subseteq \dots \subseteq F_5 \subseteq F_3 \subseteq F_1.$$

Define $H : \text{Pow}(M) \rightarrow \text{Pow}(M)$ by $H(X) := F(F(X))$. Show that H is monotone, that $\bigcup_{n \geq 0} F_{2 \cdot n}$ is the least fixed-point H_∞ of H and $\bigcap_{n \geq 0} F_{2 \cdot n+1}$ is the greatest fixed-point H_g of H . Conclude:

- F_∞ exists iff $H_\infty = H_g$;
- if F_∞ exists then $F_\infty = H_\infty$ and F_∞ is the unique fixed-point of H (and hence of F).

(b) For a concrete example of an antitone operation consider a digraph $\mathcal{G} = (G, E^G)$. The *game associated with \mathcal{G}* is played between two players I and II. The game starts in an arbitrary point of G . It is played in rounds. Each *round* consists of two moves, a move of I followed by a move of II. A player can *move* from point a to point b if $E^G ab$. A player *loses* if he cannot move. A point $a \in G$ is *won* for a player, if he has a winning strategy for games started at a . The point a is *drawn*, if for no player it is a won point. Define $F : \text{Pow}(G) \rightarrow \text{Pow}(G)$ by

$$F(X) := \{a \in G \mid \exists b \in G \setminus X : E^G ab\}.$$

Show that F is antitone and that for $n \geq 0$

- $a \in F_{2 \cdot n}$ iff a is won for I in $\leq n$ rounds;
- $a \notin F_{2 \cdot n+1}$ iff a is won for II in $\leq n + 1$ rounds.

Conclude:

- $\bigcup_{n \geq 0} F_{2 \cdot n}$ is the set of points won for I;
- $(\bigcap_{n \geq 0} F_{2 \cdot n+1})^c$ is the set of points won for II;
- F_∞ exists iff there are no drawn points;
- If F_∞ exists then F_∞ is the set of points won for I. □

A formula $\varphi(x_1, \dots, x_k, \bar{u}, X, \bar{Y})$ in the vocabulary τ , where the relation variable X has arity k , together with a τ -structure \mathcal{A} and interpretations \bar{b} and \bar{S} of \bar{u} and \bar{Y} , respectively, gives rise to an operation $F^\varphi : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$ defined by

$$F^\varphi(R) := \{(a_1, \dots, a_k) \mid \mathcal{A} \models \varphi[a_1, \dots, a_k, \bar{b}, R, \bar{S}]\}.$$

(The notation F^φ does not make explicit all relevant data.) At the beginning of this section we have already given a concrete example using this notation. A further example: For the function F introduced in part (b) of the preceding exercise for the game associated with a digraph \mathcal{G} we have $F = F^\varphi$ where

$$\varphi(x, X) := \exists y(\neg Xy \wedge E^G xy).$$

And for the function $(F^\varphi)'$ associated with F^φ as in part (c) of the preceding lemma we have

$$(+) \quad (F^\varphi)' = F^{(X \bar{x} \vee \varphi)}.$$

We now introduce *inflationary fixed-point logic* FO(IFP). It contains first-order logic and is closed under fixed-points of definable inflationary operations and hence, by 8.1.2(c), under fixed-points of definable inductive operations. Therefore, inflationary fixed-point logic is sometimes called *inductive fixed-point logic*.

The syntax of FO(IFP): For a vocabulary τ the class of formulas of FO(IFP) of vocabulary τ is given by the calculus

- $\frac{}{\varphi}$ φ an atomic second-order formula over τ
- $\frac{\varphi}{\neg\varphi}$; $\frac{\varphi, \psi}{(\varphi \vee \psi)}$; $\frac{\varphi}{\exists x\varphi}$
- $\frac{\varphi}{[\text{IFP}_{\bar{x}, X}\varphi]\bar{t}}$ where the lengths of \bar{x} and \bar{t} are the same and coincide with the arity of X .

Sentences are formulas without free first-order and second-order variables, where the free occurrence of variables is defined in the standard way, adding the clause

$$\text{free}([\text{IFP}_{\bar{x}, X}\varphi]\bar{t}) := \text{free}(\bar{t}) \cup (\text{free}(\varphi) \setminus \{\bar{x}, X\}).$$

The semantics is defined inductively w.r.t. the calculus above, the meaning of $[\text{IFP}_{\bar{x}, X}\varphi]\bar{t}$ being $\bar{t} \in F_\infty^{(X\bar{x}\vee\varphi)}$. More precisely: If X is k -ary and the variables free in $[\text{IFP}_{\bar{x}, X}\varphi]\bar{t}$ are among \bar{u} and \bar{Y} , and \bar{b} and \bar{S} are interpretations in \mathcal{A} of \bar{u} and \bar{Y} , respectively, then

$$\mathcal{A} \models [\text{IFP}_{\bar{x}, X}\varphi]\bar{t} [\bar{b}, \bar{S}] \quad \text{iff} \quad (t_1[\bar{b}], \dots, t_k[\bar{b}]) \in F_\infty^{(X\bar{x}\vee\varphi)}.$$

We sometimes denote $F_\infty^{(X\bar{x}\vee\varphi)}$ by $[\text{IFP}_{\bar{x}, X}\varphi]$.

As an example, using the formula $\chi(x, y, X) = (x = y \vee \exists z(Xxz \wedge Ezy))$ of the beginning of this section, we see that the FO(IFP)-formula in the language of graphs

$$\psi_0(x, y) := [\text{IFP}_{xy, X}(x = y \vee \exists z(Xxz \wedge Ezy))]xy$$

expresses “ $x = y$ or x, y are connected by a path”. Hence, the class of connected graphs is axiomatized by (the graph axioms and) $\forall x\forall y\psi_0(x, y)$.

A formula φ of FO(IFP) is said to be first-order, if it does not contain the IFP operator, even though it may contain second-order variables. Call an FO(IFP)-formula φ *positive (negative)* in the second-order variable X , if each free occurrence of X in φ is in the scope of an even (odd) number of negation symbols. And call φ *normal*, if for every subformula of φ of the form $[\text{IFP}_{\bar{y}, Y}\psi]\bar{t}$ the formula ψ is positive in Y . A simple induction shows (see the next exercise) that for any normal formula $\varphi(\bar{x}, X, \dots)$ positive in X the operation F^φ is monotone.

Exercise 8.1.6 (a) Let $\varphi(\bar{x}, X_1, \dots, X_k, Y_1, \dots, Y_l)$ be a normal formula of FO(IFP) that is positive in X_1, \dots, X_k and negative in Y_1, \dots, Y_l . Then $\mathcal{A} \models \varphi[\bar{a}, \bar{R}, \bar{S}], R_1 \subseteq R'_1, \dots, R_k \subseteq R'_k$, and $S'_1 \subseteq S_1, \dots, S'_l \subseteq S_l$ imply $\mathcal{A} \models \varphi[\bar{a}, \bar{R}', \bar{S}']$. (Hint: Induction on φ .)

(b) The formula $\varphi(x, X) := [\text{IFP}_{x,Y} (Xx \vee (\exists y Yy \wedge \neg Yc))]$ x is positive in X , but F^φ is not monotone. Note that φ is not normal. \square

Of course, it is decidable whether a formula φ is normal and positive in X , while – even for first-order φ – it is not decidable whether F^φ is monotone. This can be seen from Trahtenbrot’s Theorem 7.2.1 and the equivalence

$$F^{\forall x(Xx \rightarrow \psi)} \text{ is monotone (in the finite) iff } \psi \text{ is valid in the finite}$$

which holds for every first-order sentence ψ .

We saw in 8.1.2(b) that monotone operations have least fixed-points. This motivates the introduction of a corresponding fixed-point logic, namely *Least Fixed-Point Logic* FO(LFP) that we get by closing first-order logic under least fixed-points of operations definable by positive formulas. From a semantic point of view it would be more natural to consider the closure under arbitrary definable monotone operations. But then, by the preceding remark, we would get a logic with an undecidable syntax. Moreover, the restriction to normal formulas does not lead to a loss of expressive power, as even FO(IFP) is not stronger than FO(LFP) (cf. 8.1.13).

The class of formulas of FO(LFP) of vocabulary τ is the fragment of FO(IFP)[τ] given by the same calculus as that for FO(IFP), the IFP-rule being restricted to positive formulas:

- $\frac{\varphi}{[\text{IFP}_{\bar{x},X}\varphi]\bar{t}}$ where φ is positive in X and the lengths of \bar{x} and \bar{t} are the same and coincide with the arity of X .

Thus the FO(LFP)-formulas are the normal FO(IFP)-formulas. Since for normal formulas φ that are positive in X , the operation F^φ is monotone, we have, by 8.1.2, that $F_\infty^{(X\bar{x}\vee\varphi)} = F_\infty^\varphi$ and that F_∞^φ is the least fixed-point of F^φ . Hence, in this case, $[\text{IFP}_{\bar{x},X}\varphi]\bar{t}$ expresses that \bar{t} is in the least fixed-point of F^φ . We therefore write $[\text{LFP}_{\bar{x},X}\varphi]\bar{t}$ instead of $[\text{IFP}_{\bar{x},X}\varphi]\bar{t}$.

Example 8.1.7 Consider trees as $\{E\}$ -structures (cf. 1A.1). The formula of FO(LFP)

$$\varphi_{<}(x, y) := [\text{LFP}_{xy,X} (Exy \vee \exists z(Xxz \wedge Ezy))]xy$$

defines a partial ordering (an irreflexive and transitive binary relation), the *induced partial ordering*. The FO(LFP)-formula $\varphi_D(x, y) :=$

$$[\text{LFP}_{xy,X} (\forall z(\neg Ezx \wedge \neg Ezy) \vee \exists u\exists v(Xuv \wedge Eux \wedge Evy))]xy$$

expresses that x and y have the same depth and therefore, the FO(LFP)-sentence

$$\forall x \forall y (\forall z (\neg E x z \wedge \neg E y z) \rightarrow \varphi_D(x, y))$$

says that all leaves have the same depth. \square

Example 8.1.8 Fix a vocabulary τ and $s \geq 1$. We consider the pebble games introduced in subsection 3.3.1. A simple induction on n shows that for any τ -structure \mathcal{A} we have

$$F_\varphi^{n+1} := \{(a_1, \dots, a_s, b_1, \dots, b_s) \mid \text{the spoiler wins } G_n^s(\mathcal{A}, \bar{a}, \mathcal{A}, \bar{b})\},$$

where

$$\varphi(\bar{x}, \bar{y}, Z) := \bigvee_{\substack{\psi \in \text{FO}^s \\ \psi \text{ atomic}}} (\psi(\bar{x}) \leftrightarrow \neg \psi(\bar{y})) \vee \bigvee_{1 \leq i \leq s} (\exists x_i \forall y_i Z \bar{x} \bar{y} \vee \exists y_i \forall x_i Z \bar{x} \bar{y}).$$

Thus, for every $\bar{a}, \bar{b} \in A^s$ we have that $\mathcal{A} \models [\text{LFP}_{\bar{x}\bar{y}, Z} \varphi] \bar{x} \bar{y}[\bar{a}, \bar{b}]$ iff the spoiler wins $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$. \square

Exercise 8.1.9 Given a graph \mathcal{G} and $k \geq 1$, consider the following cops-and-robbers game $\text{CR}^k(\mathcal{G})$ with infinitely many moves: There are k cops and one robber, each of whom after any move stands on a vertex of \mathcal{G} . In the first move every cop chooses a position on the graph and afterwards the robber chooses his position. In any subsequent move one of the cops flies in a helicopter to a new vertex. Before the helicopter actually lands the robber will run to a new vertex along a path of the graph; however, he is not permitted to run through a cop. More formally, a position of the game $\text{CR}^k(\mathcal{G})$ is a $(k+1)$ -tuple of vertices of G . In the first move the “cops” choose $a_1, \dots, a_k \in G$ and the “robber” $b \in G$. Then, (a_1, \dots, a_k, b) is the position after the first move. Suppose (e_1, \dots, e_k, f) is the position after i moves. In the $(i+1)$ th move, first the “cops” choose l with $1 \leq l \leq k$ and $e'_l \in G$, and then the robber chooses a point f' in the connected component of f in the graph induced on $G \setminus \{e_1, \dots, e_{l-1}, e_{l+1}, \dots, e_k\}$. Then $(e_1, \dots, e_{l-1}, e'_l, e_{l+1}, \dots, e_k, f')$ is the new position. The “cops” win if eventually a position (a_1, \dots, a_k, b) with $b \in \{a_1, \dots, a_k\}$ is reached. We say that a graph \mathcal{G} has *tree-width* $\leq k$ if the “cops” have a winning strategy in $\text{CR}^{k+1}(\mathcal{G})$ (usually, this concept is defined in terms of trees).

Show that the class of graphs of tree-width $\leq k$ is axiomatizable in $\text{FO}(\text{LFP})$. Hint: Consider a formula expressing in any graph \mathcal{G} that the “cops” have a winning strategy in $\text{CR}^{k+1}(\mathcal{G})$, e.g.,

$$\exists x_1 \dots \exists x_{k+1} \forall y [\text{LFP}_{\bar{x}y, Z} (\bigvee_{i=1}^{k+1} x_i = y \vee \bigvee_{l=1}^{k+1} \exists x_l \forall z (\psi_l(\bar{x}, y, z) \rightarrow Z \bar{x} z))] \bar{x} y$$

where

$$\psi_l(\bar{x}, y, z) := [\text{LFP}_{z, U} (z = y \vee \exists u (U u \wedge E u z \wedge \bigwedge_{i \neq l} z \neq x_i))] z. \quad \square$$

For convenience we adopt the following convention:

If not explicitly stated otherwise, the notation $\varphi(\bar{x}, \bar{Y})$ only means that the variables in φ that are relevant in the given situation are among \bar{x} and \bar{Y} ; there may be other variables free in φ .

Exercise 8.1.10 (a) Let $\psi(\bar{x}, X)$ be an FO(LFP)-formula that is positive in X . Show that the greatest fixed-point of F^ψ is defined by the FO(LFP)-formula $\neg[\text{LFP}_{\bar{x}, X} \neg\psi(\bar{x}, \neg X)]\bar{x}$ (here $\psi(\bar{x}, \neg X)$ is obtained from ψ by replacing subformulas $X\bar{t}$ by $\neg X\bar{t}$; hence, $\neg\psi(\bar{x}, \neg X)$ is positive in X), a formula which sometimes is abbreviated by $[\text{GFP}_{\bar{x}, X} \psi(\bar{x}, X)]\bar{x}$. Hint: Use 8.1.4.

(b) We know that with $\varphi(x, X) := \exists y(\neg Xy \wedge Exy)$ we have $F = F^\varphi$ for the antitone function F introduced for the game associated with digraphs (cf. 8.1.5(b)). Hence, for $H := F \circ F$ we have $H = F^\psi$, where

$$\psi(x, X) := \exists y(Exy \wedge \forall z(Eyz \rightarrow Xz))$$

(note that $\psi(x, X)$ is equivalent to the formula $\varphi(x, \varphi(-, X))$ obtained from φ by replacing subformulas Xy by $\varphi(y, X)$). By Exercise 8.1.5(b) the statement “there are no drawn points” is expressed by

$$\forall x([\text{LFP}_{x, X} \psi]x \leftrightarrow [\text{GFP}_{x, X} \psi]x)$$

and the statement “ x is won by Γ ” by $[\text{LFP}_{x, X} \psi]x$. □

Example 8.1.11 Consider orderings as $\{<, S, \min, \max\}$ -structures.¹ Fix $k \geq 1$. Identify an ordering \mathcal{A} with its isomorphic copy $(\{0, \dots, n-1\}, <, S, 0, n-1)$ where $n = \|A\|$. Then every k -tuple (m_1, \dots, m_k) in A^k can be considered as the n -adic representation of a natural number, namely of

$$[m_1 \dots m_k] := m_1 \cdot n^{k-1} + \dots + m_k \cdot n^0.$$

Let S^k be the “ k -adic successor relation”,

$$S^k x_1 \dots x_k y_1 \dots y_k \quad \text{iff} \quad [y_1 \dots y_k] = [x_1 \dots x_k] + 1.$$

S^k is first-order-definable by

$$\bigvee_{1 \leq i \leq k} (x_{i+1} = \dots = x_k = \max \wedge \neg x_i = \max \wedge Sx_i y_i \wedge \bigwedge_{i < j \leq k} y_j = \min \wedge \bigwedge_{1 \leq j < i} y_j = x_j).$$

Define the relations ADD^k (“addition”) and MULT^k (“multiplication”) by

$$\text{ADD}^k x_1 \dots x_k y_1 \dots y_k z_1 \dots z_k \quad \text{iff} \quad [\bar{z}] = [\bar{x}] + [\bar{y}]$$

¹ Note that $<$, being the transitive closure of S , is definable in $\text{FO(LFP)}[\{S\}]$.

and

$$\text{MULT}^k x_1 \dots x_k y_1 \dots y_k z_1 \dots z_k \quad \text{iff} \quad [\bar{z}] = [\bar{x}] \cdot [\bar{y}].$$

ADD^k and MULT^k are definable in $\text{FO}(\text{LFP})$: For ADD^k we imitate the inductive definition

$$\begin{aligned} x + 0 &= x \\ x + (y + 1) &= (x + y) + 1 \end{aligned}$$

by an application of the LFP operator and obtain:

$$[\text{LFP}_{\bar{x}\bar{y}\bar{z},X}((\bar{y} = \widetilde{\min} \wedge \bar{z} = \bar{x}) \vee \exists \bar{v} \exists \bar{w} (X \bar{x} \bar{v} \bar{w} \wedge S^k \bar{v} \bar{y} \wedge S^k \bar{w} \bar{z}))] \bar{x} \bar{y} \bar{z},$$

a formula we denote by $\varphi_+(\bar{x}, \bar{y}, \bar{z})$. For MULT^k we use the inductive definition

$$\begin{aligned} x \cdot 0 &= 0 \\ x \cdot (y + 1) &= x \cdot y + x \end{aligned}$$

and obtain

$$[\text{LFP}_{\bar{x}\bar{y}\bar{z},X}(\bar{y} = \bar{z} = \widetilde{\min} \vee \exists \bar{v} \exists \bar{w} (X \bar{x} \bar{v} \bar{w} \wedge S^k \bar{v} \bar{y} \wedge \varphi_+(\bar{w}, \bar{x}, \bar{z})))] \bar{x} \bar{y} \bar{z},$$

a formula with “nested” LFP operators.

Now one easily sees that the graph of any polynomial is definable in $\text{FO}(\text{LFP})$ in the following sense: Given $p(x) \in \mathbb{N}[x]$, there is a k and a formula $\varphi_p(x, y_1, \dots, y_k)$ of $\text{FO}(\text{LFP})$ such that for all orderings \mathcal{A} containing the coefficients of p we have $\mathcal{A} \models \forall x \exists \bar{y} \varphi_p(x, \bar{y})$, and for all $a, \bar{b} \in A$,

$$\mathcal{A} \models \varphi_p[a, \bar{b}] \quad \text{iff} \quad p(a) = \bar{b}. \quad \square$$

The following lemma is immediate. We often use it without mentioning it explicitly. Part (b) shows how parameters can be incorporated into fixed-points.

Lemma 8.1.12 (a) *If X does not occur free in φ then $[\text{IFP}_{\bar{x},X} \varphi(\bar{x})] \bar{t}$ and $\varphi(\bar{t})$ are equivalent.*

(b) *The formulas*

$$[\text{IFP}_{\bar{x},X} \varphi(\bar{x}, y, X)] \bar{t} \quad \text{and} \quad [\text{IFP}_{\bar{x},Y} \varphi(\bar{x}, u, Y - u)] \bar{t} y$$

*are equivalent.*² □

By definition, $\text{FO}(\text{LFP}) \leq \text{FO}(\text{IFP})$. As already remarked in the introduction of this chapter we are going to show the converse:

Theorem 8.1.13 *Every $\text{FO}(\text{IFP})$ -formula is equivalent to an $\text{FO}(\text{LFP})$ -formula.*

Moreover we shall prove

² Here, $\chi(Y - u)$ is obtained from $\chi(X)$ by replacing each subformula $X\bar{s}$ by $Y\bar{s}u$.

Theorem 8.1.14 *Every FO(LFP)-formula is equivalent to an FO(LFP)-formula with at most one occurrence of LFP.*

The proofs will be given only in the next section. With the following considerations we point at an essential step.

Let X_0, X_1, \dots be the sequence of sets obtained when evaluating the outermost LFP operation in a formula of the form

$$[\text{LFP}_{\bar{x}, X} \dots [\text{LFP}_{\bar{y}, Y} \varphi(\bar{y}, X, Y)] \dots] \dots$$

In the transition from X_n to X_{n+1} one needs the value $[\text{LFP}_{\bar{y}, Y} \varphi(\bar{y}, X_n, Y)]$. We show how this kind of nested fixed-points can be expressed by “simultaneous fixed-points” and how simultaneous fixed-points can be expressed by a single fixed-point. We give the results in a general, more abstract form that will also be useful in later sections.

First, we introduce the notion of simultaneous fixed-point. Let $m \geq 0$ and suppose that M_0, \dots, M_m are finite sets and that

$$\begin{aligned}
 & F^0 : \text{Pow}(M_0) \times \dots \times \text{Pow}(M_m) \rightarrow \text{Pow}(M_0) \\
 & F^1 : \text{Pow}(M_0) \times \dots \times \text{Pow}(M_m) \rightarrow \text{Pow}(M_1) \\
 & \vdots \\
 & F^m : \text{Pow}(M_0) \times \dots \times \text{Pow}(M_m) \rightarrow \text{Pow}(M_m).
 \end{aligned}
 \tag{*}$$

We define the sequence $(F_{(n)}^0, \dots, F_{(n)}^m)_{n \geq 0}$ by

$$F_{(0)}^i := \emptyset, \quad F_{(n+1)}^i := F^i(F_{(n)}^0, \dots, F_{(n)}^m).$$

If we have for some n that $(F_{(n)}^0, \dots, F_{(n)}^m) = (F_{(n+1)}^0, \dots, F_{(n+1)}^m)$, we set

$$(F_{(\infty)}^0, \dots, F_{(\infty)}^m) := (F_{(n)}^0, \dots, F_{(n)}^m)$$

and say that the *simultaneous fixed-point* $(F_{(\infty)}^0, \dots, F_{(\infty)}^m)$ of (F^0, \dots, F^m) exists. Note that for $i = 0, \dots, m$,

$$F^i(F_{(\infty)}^0, \dots, F_{(\infty)}^m) = F_{(\infty)}^i.$$

If the simultaneous fixed-point does not exist, we set $F_{(\infty)}^i := \emptyset$ for $i \leq m$.

Example 8.1.15 (a) Let $F, G : \text{Pow}(M) \times \text{Pow}(M) \rightarrow \text{Pow}(M)$ be given by

$$F(X, Y) := M, \quad G(X, Y) := (M \setminus Y) \cap X.$$

Then $(F_{(n)})_{n \geq 0}$ is the sequence \emptyset, M, M, \dots and $(G_{(n)})_{n \geq 0}$ is the sequence $\emptyset, \emptyset, M, \emptyset, M, \emptyset, \dots$. Therefore, the simultaneous fixed-point of (F, G) does not exist and hence, $F_{(\infty)} = G_{(\infty)} = \emptyset$, even though $F_{(1)} = F_{(2)}$ and $G_{(0)} = G_{(1)}$.

(b) Suppose that the fixed-points F_{∞} and G_{∞} of $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$

and $G : \text{Pow}(N) \rightarrow \text{Pow}(N)$ exist and define the functions $\hat{F} : \text{Pow}(M) \times \text{Pow}(N) \rightarrow \text{Pow}(M)$ and $\hat{G} : \text{Pow}(M) \times \text{Pow}(N) \rightarrow \text{Pow}(N)$ by $\hat{F}(X, Y) = F(X)$ and $\hat{G}(X, Y) = G(Y)$. Then the simultaneous fixed-point $(\hat{F}_{(\infty)}, \hat{G}_{(\infty)})$ exists and $(\hat{F}_{(\infty)}, \hat{G}_{(\infty)}) = (F_{\infty}, G_{\infty})$. \square

Definition 8.1.16 Let F^0, \dots, F^m be as in (*).

(a) (F^0, \dots, F^m) is *inductive* if for $i \leq m$

$$F^i_{(0)} \subseteq F^i_{(1)} \subseteq F^i_{(2)} \subseteq \dots$$

(b) (F^0, \dots, F^m) is *inflationary* if for arbitrary X_0, \dots, X_m and all $i \leq m$,

$$X_i \subseteq F^i(X_0, \dots, X_m).$$

(c) (F^0, \dots, F^m) is *monotone* if

$$X_0 \subseteq Y_0, \dots, X_m \subseteq Y_m \text{ imply } F^i(X_0, \dots, X_m) \subseteq F^i(Y_0, \dots, Y_m)$$

for arbitrary $X_0, \dots, X_m, Y_0, \dots, Y_m$ and all $i \leq m$. \square

The proof of the following lemma is a straightforward generalization of the proof of the corresponding parts in 8.1.1 and 8.1.2.

Lemma 8.1.17 Let F^0, \dots, F^m be as in (*).

(a) If (F^0, \dots, F^m) is *inductive* then the simultaneous fixed-point exists.

(b) If (F^0, \dots, F^m) is *inflationary* or *monotone* then (F^0, \dots, F^m) is *inductive*.

(c) If (F^0, \dots, F^m) is *monotone* then $(F^0_{(\infty)}, \dots, F^m_{(\infty)})$ is the simultaneous least fixed-point of (F^0, \dots, F^m) . Even $F^i(X_0, \dots, X_m) \subseteq X_i$ for $i \leq m$ implies $F^0_{(\infty)} \subseteq X_0, \dots, F^m_{(\infty)} \subseteq X_m$.

(d) If for $i \leq m$, $G^i : \text{Pow}(M_0) \times \dots \times \text{Pow}(M_m) \rightarrow \text{Pow}(M_i)$ is given by

$$G^i(X_0, \dots, X_m) := X_i \cup F^i(X_0, \dots, X_m)$$

then (G^0, \dots, G^m) is *inflationary*. In case (F^0, \dots, F^m) is *inductive* we have $(G^0_{(n)}, \dots, G^m_{(n)}) = (F^0_{(n)}, \dots, F^m_{(n)})$ for all $n \geq 0$ and hence, $(G^0_{(\infty)}, \dots, G^m_{(\infty)}) = (F^0_{(\infty)}, \dots, F^m_{(\infty)})$. \square

We introduce extensions of FO(IFP) and FO(LFP) which allow to speak about simultaneous fixed-points. On the one hand these extensions are quite useful when formalizing statements which involve fixed-point operations. On the other hand we shall show that they do not increase the expressive power.

The formulas of *Simultaneous Inflationary Fixed-Point Logic* FO(S-IFP) are obtained by replacing the clause corresponding to IFP in the definition of FO(IFP) by

$$(+) \quad \frac{\varphi_0, \dots, \varphi_m}{[\text{S-IFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}}$$

where $m \geq 0$, where for $i = 0, \dots, m$ the arity of X_i equals the length of \bar{x}_i , and where \bar{t} has the same length as \bar{x}_0 . As usual, the variables in each sequence \bar{x}_i are distinct, but the same variable may occur in distinct \bar{x}_i 's. A variable x is free in $[\text{S-IFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$, if x is in \bar{t} or if for at least one i , x is free in φ_i and not in \bar{x}_i . We define the corresponding inflationary system (F^0, \dots, F^m) by

$$F^i(X_0, \dots, X_m) := X_i \cup \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\}$$

and let $(F^0_{(\infty)}, \dots, F^m_{(\infty)})$ be their simultaneous fixed-point. Then, by definition, $[\text{S-IFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$ means $\bar{t} \in F^0_{(\infty)}$.

Simultaneous Least Fixed-Point Logic FO(S-LFP) is defined analogously, clause (+) being restricted to formulas $\varphi_0, \dots, \varphi_m$ which are positive in all the variables X_0, \dots, X_m . The system (F^0, \dots, F^m) of corresponding functions is now defined by

$$F^i(X_0, \dots, X_m) := \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\}.$$

By positivity, the system (F^0, \dots, F^m) is monotone; hence, by 8.1.17(b),(d), $[\text{S-IFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$ expresses $\bar{t} \in F^0_{(\infty)}$. Since, by 8.1.17(c), $(F^0_{(\infty)}, \dots, F^m_{(\infty)})$ is the least fixed-point of (F^0, \dots, F^m) , we denote the formula in (+) by $[\text{S-LFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$.³

Example 8.1.18 We consider digraphs. Let F and G be the functions corresponding to the FO(S-LFP)-formulas φ_0 and φ_1 that are positive in X (unary) and Y (binary),

$$\begin{aligned} \varphi_0(x, X, Y) &:= Yxx \\ \varphi_1(x, y, X, Y) &:= Exy \vee \exists z(Yxz \wedge Ezy). \end{aligned}$$

Then $G_{(\infty)}$ is the transitive closure of the edge relation and $F_{(\infty)}$ the set of vertices z such that there is a cycle through z . Hence,

$$[\text{S-LFP}_{x, X, xy, Y} \varphi_0, \varphi_1]z$$

expresses that there is a cycle through z . Clearly, $[\text{LFP}_{xy, X} (Exy \vee \exists z(Yxz \wedge Ezy))]z$ is equivalent to this formula. Note that there is no first-order formula $\psi(x, X)$ positive in X with free variables among x, X such that $[\text{LFP}_{x, X} \psi]z$ expresses that there is a cycle through z . Otherwise, the formula $\exists z[\text{LFP}_{x, X} \psi]z$ expresses that there is a cycle. Now, $\exists z[\text{LFP}_{x, X} \psi]z$ is

³ Even though the system (F^0, \dots, F^m) corresponding to an S-IFP operator and to an S-LFP operator have different definitions, they coincide on the parts relevant for our discussion, namely on the stages and on the fixed-points.

equivalent to $F_1^\psi \neq \emptyset$ and hence, to the first-order formula $\exists x\psi(x, \emptyset)$.⁴ But we know that there is no first-order formula expressing that there is a cycle (cf. 2.3.9). \square

Notes 8.1.19 Inflationary fixed-point logic goes back to Gurevich [70]. For least fixed-point logic cf. the notes to the next section. Theorem 8.1.13 is due to Gurevich and Shelah [76] and Theorem 8.1.14 to Immerman [91]. A further reference is [18].

8.2 Simultaneous Induction and Transitivity

In the following, A will always denote a finite set with at least two elements. For $b \in A$ and $i \geq 0$ let

$$b^i := \underbrace{b \dots b}_i$$

and let \tilde{b} be b^i for some i . The length of \tilde{b} will be irrelevant or clear from the context. If Z is an $(k+l)$ -ary relation on A and $\bar{a} \in A^l$, then the \bar{a} -section $Z_{-\bar{a}}$ of Z is given by

$$Z_{-\bar{a}} := \{\bar{b} \in A^k \mid Z\bar{b}\bar{a}\}.$$

Let $m \geq 0$ and

$$(1) \quad \begin{aligned} F^0 &: \text{Pow}(A^{k_0}) \times \dots \times \text{Pow}(A^{k_m}) \rightarrow \text{Pow}(A^{k_0}) \\ F^1 &: \text{Pow}(A^{k_0}) \times \dots \times \text{Pow}(A^{k_m}) \rightarrow \text{Pow}(A^{k_1}) \\ &\quad \vdots \\ F^m &: \text{Pow}(A^{k_0}) \times \dots \times \text{Pow}(A^{k_m}) \rightarrow \text{Pow}(A^{k_m}). \end{aligned}$$

We are going to code these operations as sections of a single operation J of higher arity. Setting

$$(2) \quad k := \max\{k_0, \dots, k_m\} + (m + 1)$$

we define

$$J : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$$

in such a way that if J_0, J_1, \dots are the stages of J then for any $a, b \in A$ with $a \neq b$, the $\tilde{a}\tilde{b}^i$ -section of J_n codes $F_{(n)}^i$, that is,

$$F_{(n)}^i = J_n_{-\tilde{a}\tilde{b}^i}.$$

⁴ The formula $\psi(x, \emptyset)$ is obtained from $\psi(x, X)$ by replacing each occurrence of an atomic formula Xt by $\neg t = t$.

⁵ Hence, \tilde{a} has length $k - i - k_i$.

Lemma 8.2.1 (Simultaneous Induction) *Let F^0, \dots, F^m and k be as in (1) and (2) above. Define the function $J : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$, the simultaneous join of (F^0, \dots, F^m) , by*

$$J(Z) := \bigcup_{a,b \in A, a \neq b} ((F^0(Z_{-}\tilde{a}b^0, \dots, Z_{-}\tilde{a}b^m) \times \{\tilde{a}b^0\}) \cup \dots \cup (F^m(Z_{-}\tilde{a}b^0, \dots, Z_{-}\tilde{a}b^m) \times \{\tilde{a}b^m\})).$$

(a) *For all $n \geq 0$,*

$$J_n = \bigcup_{a,b \in A, a \neq b} ((F_{(n)}^0 \times \{\tilde{a}b^0\}) \cup \dots \cup (F_{(n)}^m \times \{\tilde{a}b^m\})).$$

Therefore, the fixed-point J_∞ of J exists iff the simultaneous fixed-point $(F_{(\infty)}^0, \dots, F_{(\infty)}^m)$ exists. Moreover,

$$J_\infty = \bigcup_{a,b \in A, a \neq b} ((F_{(\infty)}^0 \times \{\tilde{a}b^0\}) \cup \dots \cup (F_{(\infty)}^m \times \{\tilde{a}b^m\})).$$

Thus, for $i = 0, \dots, m$ and for all $a, b \in A$ with $a \neq b$,

$$F_{(n)}^i = J_n \text{ }_{-}\tilde{a}b^i \quad \text{and} \quad F_{(\infty)}^i = J_\infty \text{ }_{-}\tilde{a}b^i;$$

hence, for $\bar{c} \in A^{k^i}$,

$$F_{(\infty)}^i \bar{c} \quad \text{iff} \quad J_\infty \bar{c} \tilde{a}b^i;$$

in particular, for $\bar{c} \in A^{k^0}$,

$$(*) \quad F_{(\infty)}^0 \bar{c} \quad \text{iff} \quad J_\infty \bar{c} \tilde{a}.$$

(b) *If (F^0, \dots, F^m) is inductive then J is inductive.*

Proof. We prove by induction on n that

$$J_n = \bigcup_{a,b \in A, a \neq b} ((F_{(n)}^0 \times \{\tilde{a}b^0\}) \cup \dots \cup (F_{(n)}^m \times \{\tilde{a}b^m\}))$$

(all the other claims in (a) and (b) are then immediate). The case $n = 0$ is clear. In the induction step we have J_{n+1}

$$\begin{aligned} &= J(J_n) = J(\bigcup_{a \neq b} ((F_{(n)}^0 \times \{\tilde{a}b^0\}) \cup \dots \cup (F_{(n)}^m \times \{\tilde{a}b^m\}))) \\ &= \bigcup_{a \neq b} ((F^0(F_{(n)}^0, \dots, F_{(n)}^m) \times \{\tilde{a}b^0\}) \\ &\quad \cup \dots \cup (F^m(F_{(n)}^0, \dots, F_{(n)}^m) \times \{\tilde{a}b^m\})) \\ &= \bigcup_{a \neq b} ((F_{(n+1)}^0 \times \{\tilde{a}b^0\}) \cup \dots \cup (F_{(n+1)}^m \times \{\tilde{a}b^m\})). \end{aligned}$$

□

We come back to formulas and use the preceding lemma to prove that the extensions of FO(IFP) and FO(LFP) to the corresponding simultaneous fixed-point logics do not increase the expressive power.

To simplify the presentation we consider only structures with at least two elements; in particular, we say that φ and ψ are equivalent and write

$$\models_{\text{fin}} \varphi \leftrightarrow \psi$$

if φ and ψ are equivalent in all structures with at least two elements. After 8.2.11 we illustrate how this restriction can be removed.

First some notations that will be used frequently in the next chapters. We write

$$\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)x\psi(x),$$

if φ is equivalent to both $\exists x\psi(x)$ and $\forall x\psi(x)$. Then, for any term t , φ is equivalent to $\psi(t)$, a fact we often shall use tacitly.

To speak about distinct sections of a simultaneous fixed-point process, we use, for $l \geq 1$ and $i = 0, \dots, l$, the formulas $\delta_i^l(x_1, \dots, x_l, v, w)$ given by

$$\delta_0^l(x_1, \dots, x_l, v, w) := \neg v = w \wedge x_1 = \dots = x_l = v$$

and for $i = 1, \dots, l$ by

$$\delta_i^l(x_1, \dots, x_l, v, w) := \neg v = w \wedge x_1 = \dots = x_{l-i} = v \wedge x_{l-i+1} = \dots = x_l = w.$$

We often omit the superscript l and write $\delta_i(\bar{x}, u, v)$, l then being determined by the length of \bar{x} . Note that for distinct a, b in a structure \mathcal{A} , we have

$$\mathcal{A} \models \delta_i[\bar{a}b^j ab] \quad \text{iff} \quad i = j.$$

Furthermore, we use some more or less self-explanatory notations. For instance, given a formula $\varphi(X)$, we write $\varphi(Y\bar{y})$ for the formula obtained from φ by replacing each subformula $X\bar{t}$ by $Y\bar{t}\bar{y}$; of course, this presupposes that the arities of X and Y and the length of \bar{y} match in the right way.

Theorem 8.2.2 FO(S-LFP) \equiv FO(LFP) and FO(S-IFP) \equiv FO(IFP).

Clearly, it suffices to show that FO(S-LFP) \leq FO(LFP) and FO(S-IFP) \leq FO(IFP). The proof is by induction on formulas, the only nontrivial case being handled by the next lemma.

Lemma 8.2.3 (Simultaneous Induction for LFP and IFP) *Let*

$$\varphi_0(\bar{x}_0, X_0, \dots, X_m), \dots, \varphi_m(\bar{x}_m, X_0, \dots, X_m)$$

be formulas of FO(LFP), where X_i is k_i -ary and the length of \bar{x}_i is k_i . Furthermore, assume that $\varphi_0, \dots, \varphi_m$ are all positive in X_0, \dots, X_m . Set

$k := \max\{k_0, \dots, k_m\} + (m+1)$ and, using new variables v and w , define the FO(LFP)-formula $\chi_J(z_1, \dots, z_k, Z)$ by

$$\begin{aligned} \chi_J(z_1, \dots, z_k, Z) &:= \exists v \exists w (\neg v = w \wedge \\ &((\varphi_0(z_1, \dots, z_{k_0}, Z_- \tilde{v}w^0, \dots, Z_- \tilde{v}w^m) \wedge \delta_0(z_{k_0+1}, \dots, z_k, v, w)) \\ &\vee (\varphi_1(z_1, \dots, z_{k_1}, Z_- \tilde{v}w^0, \dots, Z_- \tilde{v}w^m) \wedge \delta_1(z_{k_1+1}, \dots, z_k, v, w)) \\ &\quad \vdots \\ &\vee (\varphi_m(z_1, \dots, z_{k_m}, Z_- \tilde{v}w^0, \dots, Z_- \tilde{v}w^m) \wedge \delta_m(z_{k_m+1}, \dots, z_k, v, w))))). \end{aligned}$$

Then, for any new variable u ,

$$\models_{\text{fin}} [\text{S-LFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t} \leftrightarrow \exists (\forall) u [\text{LFP}_{\bar{z}, Z} \chi_J] \bar{t} \tilde{u}.$$

If $\varphi_0, \dots, \varphi_m$ are first-order (existential) formulas then χ_J is first-order (existential), too. Moreover,

$$\text{free}(\chi_J) \subseteq \{\bar{z}\} \cup (\text{free}(\varphi_0) \setminus \{\bar{x}_0\}) \cup \dots \cup (\text{free}(\varphi_m) \setminus \{\bar{x}_m\}).$$

Without the assumption that $\varphi_0, \dots, \varphi_m$ are all positive in X_0, \dots, X_m the claims are true if LFP is replaced by IFP everywhere.

Proof. First, we turn to the LFP-case. For $i = 0, \dots, m$ define

$$F^i(X_0, \dots, X_m) := \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\}.$$

Then, χ_J (which is positive in Z) defines the simultaneous join of (F^0, \dots, F^m) as introduced in Lemma 8.2.1. Since $[\text{S-LFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$ expresses $F_{(\infty)}^0 \bar{t}$ and $[\text{LFP}_{\bar{z}, Z} \chi_J] \bar{t} \tilde{u}$ expresses $J_{\infty} \bar{t} \tilde{u}$, the claim follows immediately from (*) in part (a) of 8.2.1.

In the the IFP-case, for $i = 0, \dots, m$ we define F^i by

$$F^i(X_0, \dots, X_m) := X_i \cup \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\}.$$

Let J be the simultaneous join of (F^0, \dots, F^m) (cf. 8.2.1(a)). One easily verifies that $F^{(Z \bar{z} \vee \chi_J)}$ defines J' where $J'(Z) := Z \cup J(Z)$. Recall that $[\text{IFP}_{\bar{z}, Z} \chi_J] \bar{t} \tilde{u}$ means $J'_{\infty} \bar{t} \tilde{u}$. Since (F^0, \dots, F^m) is inflationary and hence, inductive, J is inductive (cf. 8.2.1(b)). Therefore, $J'_{\infty} = J_{\infty}$ (cf. 8.1.2(c)). Thus $[\text{IFP}_{\bar{z}, Z} \chi_J] \bar{t} \tilde{u}$ expresses $J_{\infty} \bar{t} \tilde{u}$ which, by (*) in 8.2.1(a), is equivalent to $F_{(\infty)}^0 \bar{t}$, that is, to

$$[\text{S-IFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}. \quad \square$$

As a first application we present a lemma which will frequently be used in the following.

Lemma 8.2.4 *Every FO(LFP)-formula $[\text{LFP}_{\bar{y},Y} \varphi] \bar{t}$ with first-order φ is equivalent to a formula of the form*

$$\exists(\forall)u[\text{LFP}_{\bar{z},Z} \psi] \bar{u}$$

with first-order ψ . If φ is existential then ψ can be chosen existential, too. The claim and the following proof remain true if LFP is replaced by IFP everywhere.

As the proof will show, we may require in addition – and will often do so tacitly – that u does not occur free in ψ and that the free variables of $\exists(\forall)u[\text{LFP}_{\bar{z},Z} \psi] \bar{u}$ are among the free variables in $[\text{LFP}_{\bar{y},Y} \varphi] \bar{t}$.

Proof. Set $\varphi_0(x, X, Y) := Y\bar{t}$ with dummy variables x and X . We show, for a new variable v ,

$$(+) \quad \models_{\text{fin}} [\text{LFP}_{\bar{y},Y} \varphi] \bar{t} \leftrightarrow [\text{S-LFP}_{x,X,\bar{y},Y} \varphi_0, \varphi] v.$$

This yields the claim: By the Lemma on Simultaneous Induction for LFP the formula on the right side is equivalent to one of the form $\exists(\forall)u[\text{LFP}_{\bar{z},Z} \psi] v\bar{u}$ with first-order ψ and $v \notin \text{free}(\psi)$ (moreover, ψ is existential if φ is existential). Hence,

$$\models_{\text{fin}} [\text{LFP}_{\bar{y},Y} \varphi] \bar{t} \leftrightarrow \exists(\forall)u[\text{LFP}_{\bar{z},Z} \psi] v\bar{u}$$

and therefore, as $v \notin \text{free}([\text{LFP}_{\bar{y},Y} \varphi] \bar{t})$,

$$\models_{\text{fin}} [\text{LFP}_{\bar{y},Y} \varphi] \bar{t} \leftrightarrow \exists(\forall)u[\text{LFP}_{\bar{z},Z} \psi] \bar{u}.$$

To prove (+), consider the functions F and G corresponding to φ_0 and φ , respectively,

$$F(X, Y) = \{x \mid \varphi_0(x, X, Y)\} \quad \text{and} \quad G(X, Y) = \{\bar{y} \mid \varphi(\bar{y}, Y)\}.$$

For the simultaneous fixed-point $(F_{(\infty)}, G_{(\infty)})$ we immediately get

$$G_{(\infty)} = F_{(\infty)}^\varphi \quad \text{and} \quad F_{(\infty)} = \begin{cases} \text{“universe”} & \text{if } G_{(\infty)} \bar{t} \\ \emptyset & \text{otherwise.} \end{cases}$$

Hence for any v ,

$$\bar{t} \in F_{(\infty)}^\varphi \quad \text{iff} \quad v \in F_{(\infty)},$$

which is a reformulation of (+). \square

One of our aims is to prove that any formula of FO(LFP) is equivalent to a formula containing only one LFP operation. As an essential step we have to express nested fixed-points by single ones. To do so we need a further result, the so-called Transitivity Lemma. It shows how the fixed-point of nested *monotone* operations can be expressed by a simultaneous fixed-point.

Lemma 8.2.5 (Transitivity) *Let (F, G) be monotone, where*

$$\begin{aligned} F &: \text{Pow}(A^k) \times \text{Pow}(A^l) \rightarrow \text{Pow}(A^k) \\ G &: \text{Pow}(A^k) \times \text{Pow}(A^l) \rightarrow \text{Pow}(A^l) \end{aligned}$$

and let $E : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$ be given by

$$E(X) := F(X, G(X, -)_\infty),$$

where $G(X, -) : \text{Pow}(A^l) \rightarrow \text{Pow}(A^l)$ denotes the monotone operation $Y \mapsto G(X, Y)$ and $G(X, -)_\infty$ denotes its least fixed-point. Then E is monotone and $E_\infty = F_{(\infty)}$.

Proof. The monotonicity of E is clear. To prove $E_\infty = F_{(\infty)}$ we show

- (1) $E(F_{(\infty)}) \subseteq F_{(\infty)}$ (hence, $E_\infty \subseteq F_{(\infty)}$ by part (b) of 8.1.2).
- (2) For all n , $F_{(n)} \subseteq E_\infty$ and $G_{(n)} \subseteq G(E_\infty, -)_\infty$ (hence, $F_{(\infty)} \subseteq E_\infty$).

For (1) we note

$$\begin{aligned} E(F_{(\infty)}) &= F(F_{(\infty)}, G(F_{(\infty)}, -)_\infty) \subseteq F(F_{(\infty)}, G_{(\infty)}) \\ &\quad (\text{since } G(F_{(\infty)}, G_{(\infty)}) = G_{(\infty)} \text{ we have } G(F_{(\infty)}, -)_\infty \subseteq G_{(\infty)}) \\ &= F_{(\infty)}. \end{aligned}$$

(2) is proved by induction: Clearly, the inclusions hold for $n = 0$. In the induction step we have

$$\begin{aligned} F_{(n+1)} &= F(F_{(n)}, G_{(n)}) \subseteq F(E_\infty, G(E_\infty, -)_\infty) \quad (\text{by induction hypothesis}) \\ &= E(E_\infty) = E_\infty \end{aligned}$$

and

$$G_{(n+1)} = G(F_{(n)}, G_{(n)}) \subseteq G(E_\infty, G(E_\infty, -)_\infty) = G(E_\infty, -)_\infty. \quad \square$$

Using the Lemma on Simultaneous Induction for LFP, the preceding lemma gives

Lemma 8.2.6 (Transitivity for LFP) *Let $\varphi(\bar{x}, X, Y)$ and $\psi(\bar{y}, X, Y)$ be first-order formulas that are positive in X and Y . Moreover, assume that no individual variable free in $[\text{LFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)]$, that is, free in ψ and distinct from \bar{y} , gets into the scope of a corresponding quantifier or LFP operator in $(*)$. Then*

$$(*) \quad [\text{LFP}_{\bar{x}, X} \varphi(\bar{x}, X, [\text{LFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)])] \bar{t}$$

is equivalent to a formula of the form

$$\exists(\forall)u[\text{LFP}_{\bar{z}, Z} \chi(\bar{z}, Z)] \tilde{u}$$

with first-order χ .

Proof. Define F and G by

$$F(X, Y) := \{\bar{x} \mid \varphi(\bar{x}, X, Y)\} \quad \text{and} \quad G(X, Y) := \{\bar{y} \mid \psi(\bar{y}, X, Y)\}.$$

Then the corresponding E in the Transitivity Lemma is given by

$$\rho_E(\bar{x}, X) := \varphi(\bar{x}, X, [\text{LFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)]).$$

Since $E_\infty = F_{(\infty)}$, we have

$$\models_{\text{fin}} [\text{LFP}_{\bar{x}, X} \varphi(\bar{x}, X, [\text{LFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)])] \bar{t} \leftrightarrow [\text{S-LFP}_{\bar{x}, X, \bar{y}, Y} \varphi, \psi] \bar{t}.$$

By 8.2.3, the formula on the right side is equivalent to $[\text{LFP}_{\bar{z}, Z} \chi_0] \bar{t} \bar{v}$ for suitable first-order χ_0 and hence, by 8.2.4, to one of the desired form. \square

In a first step towards a proof of 8.1.14 we use the Transitivity Lemma to settle the case of FO(LFP)-formulas with only positive occurrences of the LFP operator.

Lemma 8.2.7 *Every FO(LFP)-formula φ containing only positive occurrences of the LFP operator (i.e., each LFP operator in φ is in the scope of an even number of negation symbols) is equivalent to a formula of the form*

$$\exists (\forall) u [\text{LFP}_{\bar{z}, Z} \psi] \bar{u}$$

where ψ is first-order.

Proof. By 8.2.4 it suffices to give a representation of φ in the form

$$[\text{LFP}_{\bar{z}, Z} \psi] \bar{t}$$

with first-order ψ . We use the connectives \neg, \wedge, \vee and the quantifiers \forall and \exists . Then, by hypothesis, we can assume that all negation symbols in φ are in front of atomic formulas. We proceed by induction on φ .

If φ is atomic or the negation of an atomic formula then

$$\models_{\text{fin}} \varphi \leftrightarrow [\text{LFP}_{x, X} \varphi] x$$

for any dummy variables x and X .

Suppose $\varphi = (\varphi_1 \wedge \varphi_2)$. By induction hypothesis we can assume that

$$\varphi_1 = [\text{LFP}_{\bar{x}, X} \psi_1(\bar{x}, X)] \bar{s} \quad \text{and} \quad \varphi_2 = [\text{LFP}_{\bar{y}, Y} \psi_2(\bar{y}, Y)] \bar{t}$$

with first-order ψ_1 and ψ_2 . Moreover, we can assume that no variable in \bar{x} occurs in ψ_2 . Then one easily verifies that

$$\models_{\text{fin}} (\varphi_1 \wedge \varphi_2) \leftrightarrow [\text{LFP}_{\bar{x}, X} (\psi_1(\bar{x}, X) \wedge [\text{LFP}_{\bar{y}, Y} \psi_2(\bar{y}, Y)] \bar{t})] \bar{s}$$

which, by the Transitivity Lemma 8.2.6 (for $\varphi(\bar{x}, X, Y) := \psi_1(\bar{x}, X) \wedge Y\bar{t}$ and for $\psi(\bar{y}, X, Y) := \psi_2(\bar{y}, Y)$), is equivalent to a formula of the desired form.

The proof for $\varphi = (\varphi_1 \vee \varphi_2)$ is completely analogous.

Suppose that $\varphi = \forall v\psi_1$. By induction hypothesis we can assume that $\varphi = \forall v[\text{LFP}_{\bar{y}, Y} \psi_1] \bar{t}$ with a first-order formula ψ_1 and, by 8.1.12(b), that the variables free in ψ_1 are among \bar{y} . By 8.1.12(a), for dummy x and X ,

$$\models_{\text{fin}} \varphi \leftrightarrow [\text{LFP}_{x, X} \forall v[\text{LFP}_{\bar{y}, Y} \psi_1] \bar{t}] v.$$

By the Transitivity Lemma (for $\varphi(x, X, Y) := \forall v Y\bar{t}$ and $\psi(\bar{y}, X, Y) := \psi_1$), this formula is equivalent to a formula of the form claimed.

The proof for $\varphi = \exists v\psi_1$ is completely analogous.

Finally suppose that

$$\varphi = [\text{LFP}_{\bar{x}, X} \varphi_1] \bar{t}$$

where, by induction hypothesis, we may assume that $\varphi_1 = [\text{LFP}_{\bar{y}, Y} \psi_1] \bar{s}$ with first-order ψ_1 and, by 8.1.12(b), that no variable in \bar{x} is free in $[\text{LFP}_{\bar{y}, Y} \psi_1]$. Hence,

$$\varphi = [\text{LFP}_{\bar{x}, X} [\text{LFP}_{\bar{y}, Y} \psi_1] \bar{s}] \bar{t}.$$

Once more, the Transitivity Lemma (for the formulas $\varphi(\bar{x}, X, Y) := Y\bar{s}$ and $\psi(\bar{y}, X, Y) := \psi_1$) yields the desired form. \square

We now want to extend the preceding lemma to arbitrary formulas of FO(LFP). This will easily be obtained by the methods we are going to develop in order to show that every FO(IFP)-formula is equivalent to an FO(LFP)-formula. So we turn to this problem next.

Let $\varphi(\bar{x}, X)$ be an FO(LFP)-formula. Recall that $[\text{IFP}_{\bar{x}, X} \varphi(\bar{x}, X)] \bar{t}$ expresses that $\bar{t} \in F_\infty^{(X\bar{x} \vee \varphi(\bar{x}, X))}$. Note that X may occur positively *and* negatively in $(X\bar{x} \vee \varphi)$. Replace all (free) negative occurrences of X by $\neg Y$ (where Y is a new variable), thus getting an FO(LFP)-formula $\psi(\bar{x}, X, Y)$, which is positive in X and Y , such that

$$(*) \quad (X\bar{x} \vee \varphi(\bar{x}, X)) \text{ is (equivalent to) } \psi(\bar{x}, X, \neg X)$$

(for example, if $\varphi(x, y, X) = \neg \exists u (\neg Xxu \vee Xuy)$, then $\psi(x, y, X, Y) = (Xxy \vee \neg \exists u (\neg Xxu \vee \neg Yuy))$). Since ψ is positive in X and Y , the function

$$L(X, Y) := \{\bar{x} \mid \psi(\bar{x}, X, Y)\}$$

is monotone (cf. 8.1.6). We set $H(X) := L(X, X^c)$.⁶ Then $H = F^{\psi(\bar{x}, X, \neg X)} = F^{(X\bar{x} \vee \varphi(\bar{x}, X))}$ and therefore, H is inflationary. We are going to show how the fixed-point of H can be expressed in terms of two relations, the “stage

⁶ For $X \subseteq M$, X^c denotes the complement $M \setminus X$.

comparison relations". These relations can be obtained as simultaneous fixed-points of positive formulas and thus, by Simultaneous Induction for FO(LFP), as a fixed-point of a positive formula, i.e., by an FO(LFP)-formula. Once more we study parts of the problem on a more abstract level.

First we introduce the *rank function* $| \cdot |_H$ and the *stage comparison relations* \leq_H and \prec_H associated with an arbitrary inductive $H : \text{Pow}(M) \rightarrow \text{Pow}(M)$ (with finite M). We define $| \cdot |_H : M \rightarrow \mathbb{N} \cup \{\infty\}$ by

$$|a|_H := \begin{cases} n & \text{if } a \in H_n \setminus H_{n-1} \\ \infty & \text{if } a \notin H_\infty \end{cases}$$

and the binary relations \leq_H and \prec_H on M by

$$\begin{array}{ll} a \leq_H b & \text{iff } a, b \in H_\infty \text{ and } |a|_H \leq |b|_H \\ a \prec_H b & \text{iff } a \in H_\infty \text{ and } |a|_H < |b|_H, \end{array}$$

where, by definition, $n < \infty$ for $n \in \mathbb{N}$.

Parts (a) – (e) of the next lemma are immediate from the definitions.

Lemma 8.2.8 (a) $a \in H_\infty$ iff $a \leq_H a$.

(b) For $a \in H_\infty$,

$$\{u \mid u \leq_H a\}^c = \{u \mid a \prec_H u\}.$$

(c) $a \leq_H b$ iff $a, b \in H_1 \vee \exists e (e \prec_H b \wedge a, b \in H(\{u \mid u \leq_H e\}))$

(in words: $a \leq_H b$ iff both a, b are contained in the first stage H_1 or there is a stage not containing b such that a and b are members of the next stage).

(d) $a \prec_H b$ iff $(a \in H_1 \wedge b \notin H_1) \vee \exists e (e \leq_H e \wedge a \in H(\{u \mid u \leq_H e\}) \wedge b \notin H(\{u \mid u \leq_H e\}))$.

(e) If $n_0 \geq 1$ and $H_\infty = H_{n_0} \neq H_{n_0-1}$ then

$$a \notin H_\infty \quad \text{iff} \quad \exists b (b \in H_{n_0} \setminus H_{n_0-1} \wedge b \prec_H a).$$

(f) If $n_0 \geq 1$ and $H_\infty = H_{n_0} \neq H_{n_0-1}$ then

$$b \in H_{n_0} \setminus H_{n_0-1} \quad \text{iff} \quad \forall e (e \leq_H b \vee e \notin H(\{u \mid u \leq_H b\})).$$

Proof. (f) First assume that b satisfies the right side. Then, $b \in H_\infty$, as otherwise, for all $e, e \notin H(\{u \mid u \leq_H b\}) = H(\emptyset) = H_1$, hence, $H_1 = \emptyset = H_\infty$, a contradiction. If $b \in H_n \setminus H_{n-1}$ for some $n < n_0$, then no $e \in H_{n+1} \setminus H_n$ would satisfy $(e \leq_H b \vee e \notin H(\{u \mid u \leq_H b\}))$, a contradiction. Conversely, suppose $b \in H_{n_0} \setminus H_{n_0-1}$ and let e be arbitrary: if $e \in H_\infty$ then $e \leq_H b$; if $e \notin H_\infty$ then $e \notin H(\{u \mid u \leq_H b\})$. \square

Now suppose that for the inductive function H there is a monotone function $L : \text{Pow}(M) \times \text{Pow}(M) \rightarrow \text{Pow}(M)$ such that

$$\text{for all } X \subseteq M, \quad H(X) = L(X, X^c).$$

Then the operation $\widehat{L} : \text{Pow}(M) \times \text{Pow}(M) \rightarrow \text{Pow}(M)$ which is given by

$$\text{for all } X, Y \subseteq M, \quad \widehat{L}(X, Y) := (L(X^c, Y^c))^c$$

is monotone, too: If $X_1 \subseteq X_2$ and $Y_1 \subseteq Y_2$ we obtain stepwise

$$X_2^c \subseteq X_1^c, \quad Y_2^c \subseteq Y_1^c, \quad L(X_2^c, Y_2^c) \subseteq L(X_1^c, Y_1^c), \quad \widehat{L}(X_1, Y_1) \subseteq \widehat{L}(X_2, Y_2).$$

Using part (b) of the preceding lemma we can rewrite parts (c) and (d) as

$$(c') \quad a \leq_H b \quad \text{iff} \quad a, b \in H_1 \vee \exists e (e \prec_H b \wedge a, b \in L(\{u \mid u \leq_H e\}, \{u \mid e \prec_H u\}))$$

$$(d') \quad a \prec_H b \quad \text{iff} \quad (a \in H_1 \wedge b \notin H_1) \vee \exists e (e \leq_H e \wedge a \in L(\{u \mid u \leq_H e\}, \{u \mid e \prec_H u\}) \wedge b \in \widehat{L}(\{u \mid e \prec_H u\}, \{u \mid u \leq_H e\})).$$

We now define operations $F, G : \text{Pow}(M^2) \times \text{Pow}(M^2) \rightarrow \text{Pow}(M^2)$ in such a way that

$$F(\leq_H, \prec_H) = \leq_H \quad \text{and} \quad G(\leq_H, \prec_H) = \prec_H.$$

(c') and (d') show that this is achieved by setting (we use notations such as $a F(U, V) b$ for $(a, b) \in F(U, V)$)

$$(1) \quad a F(U, V) b \quad \text{iff} \quad a, b \in H_1 \vee \exists e (e V b \wedge a, b \in L(_U e, e V _))$$

$$(2) \quad a G(U, V) b \quad \text{iff} \quad (a \in H_1 \wedge b \notin H_1) \vee \exists e (e U e \wedge a \in L(_U e, e V _) \wedge b \in \widehat{L}(e V _, _U e)).$$

By monotonicity of L and \widehat{L} , the functions F and G are monotone, too. Once more, we define $F_{(n)}$ and $G_{(n)}$ by

$$F_{(0)} := \emptyset, \quad G_{(0)} := \emptyset, \quad F_{(n+1)} := F(F_{(n)}, G_{(n)}), \quad G_{(n+1)} := G(F_{(n)}, G_{(n)}).$$

Then we have

$$(3) \quad a F_{(n)} b \quad \text{iff} \quad a, b \in H_n \wedge a \leq_H b$$

$$(4) \quad a G_{(n)} b \quad \text{iff} \quad a \in H_n \wedge a \prec_H b.$$

The proof is by simultaneous induction, the induction step for (3) being:

$$\begin{aligned}
 a F_{(n+1)} b & \text{ iff } a F(F_{(n)}, G_{(n)}) b \\
 & \text{ iff } a, b \in H_1 \vee \exists e (e G_{(n)} b \wedge a, b \in L(-F_{(n)}e, eG_{(n)}-)) \\
 & \text{ iff } a, b \in H_1 \vee \exists e (e \in H_n \wedge e \prec_H b \wedge \\
 & \quad a, b \in L(\{u \mid u \leq_H e\}, \{u \mid e \prec_H u\})) \text{ (by ind. hyp.)} \\
 & \text{ iff } a, b \in H_1 \vee \exists e (e \in H_n \wedge e \prec_H b \wedge \\
 & \quad a, b \in H(\{u \mid u \leq_H e\}) \text{ (by choice of } L)) \\
 & \text{ iff } a, b \in H_{n+1} \text{ and } a \leq_H b \text{ (by definition of } H_{n+1}).
 \end{aligned}$$

From (3) and (4) we get

$$(5) \quad (F_{(\infty)}, G_{(\infty)}) = (\leq_H, \prec_H).$$

Thus (\leq_H, \prec_H) is the simultaneous fixed-point of the monotone system (F, G) .

We come back to our original aim of expressing the result of an IFP operation by an application of the LFP operator.

Suppose that $\varphi(\bar{x}, X)$ is an FO(LFP)-formula. Choose an FO(LFP)-formula $\psi(\bar{x}, X, Y)$ positive in X and Y such that

$$(X\bar{x} \vee \varphi(\bar{x}, X)) \quad \text{and} \quad \psi(\bar{x}, X, \neg X)$$

are equivalent. As (the inductive) H and (the monotone) L of the preceding discussion take the operations given by

$$H(X) := \{\bar{x} \mid X\bar{x} \vee \varphi(\bar{x}, X)\} \quad \text{and} \quad L(X, Y) := \{\bar{x} \mid \psi(\bar{x}, X, Y)\},$$

respectively. Then $H(X) = L(X, X^c)$, $\widehat{L}(X, Y) = \{\bar{x} \mid \neg\psi(\bar{x}, \neg X, \neg Y)\}$, and $\neg\psi(\bar{x}, \neg X, \neg Y)$ is positive in X and Y . The first stage H_1 of H is defined by $\varphi(\bar{x}, \emptyset)$, where $\varphi(\bar{x}, \emptyset)$ is obtained from $\varphi(\bar{x}, X)$ by replacing each occurrence of an atomic part of the form $X\bar{t}$ by $\neg t_1 = t_1$. By (1) and (2), the corresponding operations F and G can be defined by the FO(LFP)-formulas

$$\begin{aligned}
 \epsilon(\bar{x}, \bar{y}, U, V) & := (\varphi(\bar{x}, \emptyset) \wedge \varphi(\bar{y}, \emptyset)) \vee \exists \bar{z} (\bar{z} V \bar{y} \wedge \\
 & \quad \psi(\bar{x}, -U\bar{z}, \bar{z}V_-) \wedge \psi(\bar{y}, -U\bar{z}, \bar{z}V_-))
 \end{aligned}$$

$$\begin{aligned}
 \delta(\bar{x}, \bar{y}, U, V) & := (\varphi(\bar{x}, \emptyset) \wedge \neg\varphi(\bar{y}, \emptyset)) \vee \exists \bar{z} (\bar{z} U \bar{z} \wedge \\
 & \quad \psi(\bar{x}, -U\bar{z}, \bar{z}V_-) \wedge \neg\psi(\bar{y}, \neg\bar{z}V_-, \neg-U\bar{z})),
 \end{aligned}$$

respectively. Both are positive in U, V . As $(\leq_H, \prec_H) = (F_{(\infty)}, G_{(\infty)})$ by (5), the relations \leq_H and \prec_H are definable by

$$\begin{aligned}
 \varphi_{\leq_H}(\bar{x}, \bar{y}) & := [\text{S-LFP}_{\bar{x}\bar{y}, U, \bar{x}\bar{y}, V} \epsilon, \delta] \bar{x}\bar{y}, \\
 \varphi_{\prec_H}(\bar{x}, \bar{y}) & := [\text{S-LFP}_{\bar{x}\bar{y}, V, \bar{x}\bar{y}, U} \delta, \epsilon] \bar{x}\bar{y},
 \end{aligned}$$

respectively. By 8.2.3 there are FO(LFP)-formulas ψ_{\leq_H} and $\psi_{<_H}$ such that

$$\begin{aligned} \models_{\text{fin}} \varphi_{\leq_H}(\bar{x}, \bar{y}) &\leftrightarrow \exists(\forall)u[\text{LFP}_{\bar{z}, Z} \psi_{\leq_H}] \tilde{u}, \\ \models_{\text{fin}} \varphi_{<_H}(\bar{x}, \bar{y}) &\leftrightarrow \exists(\forall)u[\text{LFP}_{\bar{z}, Z} \psi_{<_H}] \tilde{u}. \end{aligned}$$

Finally, using part (a) of 8.2.8, we get that

$$\models_{\text{fin}} [\text{IFP}_{\bar{x}, X} \varphi] \bar{t} \leftrightarrow \varphi_{\leq_H}(\bar{t}, \bar{t}).$$

Altogether, we can replace an IFP operation by an equivalent LFP operation. This immediately allows us to handle the only nontrivial step of an inductive proof of Theorem 8.1.13.

With the following considerations we turn to the proof of Theorem 8.1.14. First, we show that every formula $\neg[\text{LFP}_{\bar{x}, X} \varphi(\bar{x}, X)] \bar{t}$ with first-order φ is equivalent to an FO(LFP)-formula containing only positive occurrences of the LFP operator. So let $\varphi(\bar{x}, X)$ be first-order. Then, the formulas ϵ and δ are first-order and, by 8.2.3, ψ_{\leq_H} and $\psi_{<_H}$ can be chosen first-order. If $n_0 \geq 1$ and $H_\infty = H_{n_0} \neq H_{n_0-1}$ then, by 8.2.8(f) and by the definition of \hat{L} , we have

$$b \in H_{n_0} \setminus H_{n_0-1} \text{ iff } \forall e (e \leq_H b \vee e \in \hat{L}(\{u \mid b <_H u\}, \{u \mid u \leq_H b\})).$$

Thus, $\bar{z} \in H_{n_0} \setminus H_{n_0-1}$ can be expressed by

$$\varphi_{\max}(\bar{z}) := \forall \bar{v} (\varphi_{\leq_H}(\bar{v}, \bar{z}) \vee \neg \psi(\bar{v}, \neg \varphi_{<_H}(\bar{z}, -), \neg \varphi_{\leq_H}(-, \bar{z})).$$

Together with 8.2.8(e) we get the equivalence of $\neg[\text{LFP}_{\bar{x}, X} \varphi(\bar{x}, X)] \bar{t}$ and

$$(*) \quad \forall \bar{x} \neg \varphi(\bar{x}, \emptyset) \vee (\exists \bar{x} \varphi(\bar{x}, \emptyset) \wedge \exists \bar{z} (\varphi_{\max}(\bar{z}) \wedge \varphi_{<_H}(\bar{z}, \bar{t}))).$$

Since φ_{\leq_H} and $\varphi_{<_H}$ only contain positive occurrences of the LFP operator, the same applies to the formula in (*).

Thus, every FO(LFP)-formula is equivalent to one with only positive occurrences of LFP. In view of Lemma 8.2.7 this finishes the proof of 8.1.14 and, in fact, gives the following strengthening of it:

Theorem 8.2.9 *For every FO(LFP)-formula φ there is a first-order formula ψ such that*

$$\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)u[\text{LFP}_{\bar{z}, Z} \psi] \tilde{u}. \quad \square$$

Corollary 8.2.10 *For every FO(LFP)-formula φ with free variables among $\bar{x} = x_1 \dots x_n$ with $n \geq 1$ there are first-order formulas ψ_0, ψ_1 with free individual variables among \bar{x} and \bar{y} , respectively, such that*

$$\models_{\text{fin}} \varphi(\bar{x}) \leftrightarrow [\text{S-LFP}_{\bar{x}, X, \bar{y}, Y} \psi_0, \psi_1] \bar{x}.$$

Proof. With the preceding theorem choose a first-order formula $\psi(\bar{x}, \bar{z}, Z)$ such that φ and $\exists(\forall)u[\text{LFP}_{\bar{z}, Z} \psi] \tilde{u}$ are equivalent. We can assume that $\{\bar{z}\} \cap \{\bar{x}\} = \emptyset$ and that the free individual variables of ψ are among $\bar{x} \bar{z}$. Let Y be of arity $\text{length}(\bar{x}) + \text{length}(\bar{z})$. Then

$$\begin{aligned} \psi_0(\bar{x}, X, Y) &:= Y \bar{x} \tilde{x}_1 \\ \psi_1(\bar{x}, \bar{z}, X, Y) &:= \psi(\bar{x}, \bar{z}, Y \bar{x}_-) \end{aligned}$$

satisfy the claim, where $\bar{y} = \bar{x} \bar{z}$. \square

We saw in 8.1.18 that in the preceding corollary we cannot replace the formula $[\text{S-LFP}_{\bar{x}, X, \bar{y}, Y} \psi_0, \psi_1] \bar{x}$ by a formula of the form $[\text{LFP}_{\bar{x}, X} \psi] \bar{x}$: a counterexample is given by a formula $\varphi(x)$ expressing in digraphs that there is a cycle through x .

Exercise 8.2.11 Consider a formula $\chi = \exists(\forall)u[\text{IFP}_{\bar{x}, X} \varphi] \tilde{u}$ where φ is first-order and Σ_1 (cf. page 7). Show that χ is equivalent to a formula $\exists(\forall)u[\text{LFP}_{\bar{z}, Z} \rho] \tilde{u}$ where ρ is first-order and Δ_2 (that is, equivalent to a Σ_2 -formula and to a Π_2 -formula). Hint: Assume that φ is a Σ_1 -formula and let ψ, ϵ, δ be as in the exposition following 8.2.8. Then ϵ is equivalent to a Σ_1 -formula and δ to a Σ_2 -formula; hence, the corresponding $\psi_{\leq H}$ is equivalent to a Σ_2 -formula. To obtain a Π_2 representation, note that 8.2.8(d) can be replaced by

$$\begin{aligned} a \prec_H b \quad \text{iff} \quad & (a \leq_H a \wedge b \notin H_1) \wedge \\ & \forall e(a \leq_H e \vee a \prec_H e \vee b \notin H(\{u \mid u \leq_H e\})). \end{aligned}$$

Therefore, as $\delta(\bar{x}, \bar{y}, U, V)$ we can use the Π_1 -formula

$$\bar{x} U \bar{x} \wedge \neg \varphi(\bar{y}, \emptyset) \wedge \forall \bar{z}(\bar{x} U \bar{z} \vee \bar{x} V \bar{z} \vee \neg \psi(\bar{y}, \neg \bar{z} V_-, \neg U \bar{z})).$$

Then the corresponding χ_J is equivalent to a Π_2 -formula (note that χ_J has the form $\exists v \exists w(\neg v = w \wedge \theta)$, but is equivalent to $\forall v \forall w(\neg v = w \rightarrow \theta)$). An inspection of the proof of 8.2.4 shows that we get a Δ_2 -formula if we start from a Δ_2 -formula. \square

As promised, we now show how one can remove the (tacitly assumed) restriction to structures with at least two elements. As an example we take the preceding theorem (Theorem 8.2.9). First of all we note that in structures with only one element the formulas

$$[\text{LFP}_{\bar{v}, V} \chi(\bar{v}, V)] \bar{t} \quad \text{and} \quad \chi(\bar{t}, \emptyset)$$

are equivalent. Hence, for every FO(LFP)-formula ρ there is a first-order formula ρ^* equivalent to it in these structures. Now, given $\varphi \in \text{FO(LFP)}$, we use the preceding theorem to choose a first-order ψ such that φ and $\exists(\forall)u[\text{LFP}_{\bar{z}, Z} \psi] \tilde{u}$ are equivalent in structures with at least two elements.

Then, for new x and y , the formulas φ and $\exists(\forall)u[\text{LFP}_{\bar{z},Z}(\forall x\forall yx = y \wedge \varphi^*) \vee (\exists x\exists y\neg x = y \wedge \psi)] \tilde{u}$ are equivalent in any finite structure.

We close this section by giving a further result demonstrating the value of simultaneous fixed-point logic: We show that, on ordered structures, it captures the class of polynomial time computable functions in terms of the stages of simultaneous fixed-points. First we introduce the corresponding terminology. Let

$$\varphi_0(\bar{x}_0, X_0, \dots, X_m), \dots, \varphi_m(\bar{x}_m, X_0, \dots, X_m)$$

be first-order formulas that are positive in X_0, \dots, X_m , where $\text{arity}(X_i) = \text{length}(\bar{x}_i)$ for $i \leq m$. Given a structure \mathcal{A} , let F^0, \dots, F^m be the functions corresponding to $\varphi_0, \dots, \varphi_m$,

$$F^i(X_0, \dots, X_m) = \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\}.$$

For $i \leq m$, let $|\varphi_i|(\mathcal{A})$ be the number of different stages of F^i , more precisely,

$$|\varphi_i|(\mathcal{A}) := \|\{n \mid F_{(n)}^i \neq F_{(n+1)}^i\}\|.$$

Clearly, the function $\mathcal{A} \mapsto |\varphi_i|(\mathcal{A})$ is a function computable in polynomial time. A converse is also true:

Theorem 8.2.12 *Let $< \in \tau$ and let K be a class of ordered τ -structures in PTIME. Then for every function $f : K \rightarrow \mathbb{N}$ computable in polynomial time⁷ there are first-order formulas $\varphi_0(\bar{x}_0, X_0, \dots, X_m), \dots, \varphi_m(\bar{x}_m, X_0, \dots, X_m)$ that are positive in X_0, \dots, X_m such that for all $\mathcal{A} \in K$ with at least two elements,*

$$f(\mathcal{A}) = |\varphi_0|(\mathcal{A}).$$

Proof. Since f is computable in polynomial time, there is an r such that $f(\mathcal{A}) < \|A\|^r$ for all $\mathcal{A} \in K$ with at least two elements. For $j < \|A\|^r$ let $|j|_r$ be the j -th r -tuple in the lexicographic ordering of A^r . By assumption, the class

$$\{(\mathcal{A}, a_1, \dots, a_r) \mid \mathcal{A} \in K, |f(\mathcal{A})|_r = (a_1, \dots, a_r)\}$$

is in PTIME. Hence, by 7.3.4 and 8.1.13, there is a formula $\varphi(x_1, \dots, x_r)$ of FO(LFP) such that for any ordered structure \mathcal{B} and $\bar{b} \in B^r$,

$$\mathcal{B} \models \varphi[\bar{b}] \quad \text{iff} \quad \mathcal{B} \in K \text{ and } \bar{b} = |f(\mathcal{B})|_r.$$

By 8.2.10, we obtain first-order formulas $\psi_0(\bar{x}, X, Y), \psi_1(\bar{y}, X, Y)$ that are positive in X, Y such that $\varphi(\bar{x})$ and $[\text{S-LFP}_{\bar{x}, X, \bar{y}, Y} \psi_0, \psi_1] \bar{x}$ are equivalent. To get the desired representation of $f(\mathcal{A})$, we construct a simultaneous fixed-point process, where the stages of the first component Z (of arity r) successively take up the r -tuples in lexicographic order that are smaller than $|f(\mathcal{A})|_r$ as soon as $|f(\mathcal{A})|_r$ gets into X .

⁷ We assume that $\mathcal{A} \cong \mathcal{B}$ implies $f(\mathcal{A}) = f(\mathcal{B})$.

We set

$$\psi(\bar{z}, Z, X, Y) := \exists \bar{x}(X\bar{x} \wedge \bar{z} < \bar{x} \wedge (\bar{z} = \widetilde{\min} \vee \exists \bar{u}(Z\bar{u} \wedge \bar{z} = \bar{u} + 1))).$$

For the system $\psi(\bar{z}, Z, X, Y)$, $\psi_0(\bar{x}, Z, X, Y)$, $\psi_1(\bar{y}, Z, X, Y)$ one easily verifies that $f(\mathcal{A}) = |\psi|(\mathcal{A})$. \square

In the preceding theorem we cannot replace $\varphi_0, \dots, \varphi_m$ by a single formula. To give a counterexample, we consider the function f defined on the class of ordered graphs by

$$f(\mathcal{G}) := \begin{cases} 1 & \text{if } \mathcal{G} \text{ is connected} \\ 0 & \text{otherwise.} \end{cases}$$

Suppose that, for some first-order formula $\varphi(\bar{x}, X)$ positive in X , we have $f(\mathcal{G}) = |\varphi|(\mathcal{G})$. Then, \mathcal{G} is connected iff $F_\infty^\varphi \neq \emptyset$. But $F_\infty^\varphi \neq \emptyset$ is equivalent to $F_1^\varphi \neq \emptyset$. Hence, \mathcal{G} is connected iff \mathcal{G} is a model of $\exists \bar{x}\varphi(\bar{x}, \emptyset)$. But we saw in 2.3.7 that the class of finite connected ordered graphs is not first-order axiomatizable.

Notes 8.2.13 The book [122] contains the first systematic study of least fixed-point operators (in the context of infinite structures); in particular, in this book the Lemma on Simultaneous Induction and the Transitivity Lemma are proven and the stage comparison relations are introduced. The result of Exercise 8.2.11 is due to Dahlhaus [23], Theorem 8.2.12 to Kolaitis and Thakur [103].

8.3 Partial Fixed-Point Logic

In Chapter 7 a further fixed-point logic turned out to be relevant, namely partial fixed-point logic, its importance being documented by the fact that, in ordered structures, exactly the queries computable in polynomial space are expressible in partial fixed-point logic. In this section we recall the definition of partial fixed-point logic, discuss the scope of totally defined fixed-points, and show that every formula of partial fixed-point logic is equivalent to a formula containing exactly one occurrence of the PFP operator.

Partial Fixed-Point Logic FO(PFP) contains first-order logic and “ F_∞ for any definable F ”: For a vocabulary τ the class of formulas of FO(PFP) of vocabulary τ is given by the calculus

- $\frac{}{\varphi}$ where φ is an atomic second-order formula over τ
- $\frac{\varphi}{\neg\varphi}, \frac{\varphi, \psi}{(\varphi \vee \psi)}, \frac{\varphi}{\exists x\varphi}$

- $\frac{\varphi}{[\text{PFP}_{\bar{x},X}\varphi]\bar{t}}$ where the lengths of \bar{x} and \bar{t} are the same and coincide with the arity of X .

The semantics is defined inductively w.r.t. this calculus, the meaning of $[\text{PFP}_{\bar{x},X}\varphi]\bar{t}$ being $\bar{t} \in F_\infty^\varphi$. In particular, $[\text{PFP}_{\bar{x},X}\varphi]\bar{t}$ is false if the fixed-point of F^φ does not exist. All other definitions and conventions are as for FO(LFP).

Clearly, we have

Proposition 8.3.1 $\text{FO(LFP)} \leq \text{FO(PFP)}$. □

Consider, for example, orderings as $\{<, S, \min, \max\}$ -structures. Then the sentence $\exists x [\text{PFP}_{x,X} \psi(x, X)] x$ where

$$\psi(x, X) = (\forall y \neg Xy \wedge x = \min) \vee (X \max \wedge x = \max) \vee \exists y (Xy \wedge \exists u (Syu \wedge Sux))$$

just holds in orderings of odd length.

As a further example consider the formula

$$[\text{PFP}_{x,X} \exists y (\neg Xy \wedge Exy)] x.$$

By 8.1.5(b), in digraphs it expresses, at the same time, that x is won for player I and that there are no drawn points. By 8.1.10(b) we see that this formula is equivalent to an FO(LFP)-formula. This result is generalized in the following (cf. 8.1.5(a))

Exercise 8.3.2 Assume that $\varphi(\bar{x}, X)$ is an FO(LFP)-formula negative in X and that X does not occur in the scope of any LFP operator. Then

- $\psi(\bar{x}, X) := \varphi(\bar{x}, \varphi(-, X))$ is an FO(LFP)-formula positive in X .
- F^φ is antitone and $F^\psi = F^\varphi \circ F^\varphi$.
- The “FO(LFP,PFP)”-formula $[\text{PFP}_{\bar{x},X} \varphi(\bar{x}, X)]\bar{t}$ is equivalent to the formula of FO(LFP)

$$\forall \bar{x} ([\text{LFP}_{\bar{x},X} \psi]\bar{x} \leftrightarrow [\text{GFP}_{\bar{x},X} \psi]\bar{x}) \wedge [\text{LFP}_{\bar{x},X} \psi]\bar{t}. \quad \square$$

We aim at showing

Theorem 8.3.3 *Every FO(PFP)-formula is equivalent to an FO(PFP)-formula with at most one occurrence of PFP.*

We shall point out similarities and differences to the proof of the corresponding result for FO(LFP). As there, we start by introducing *Simultaneous Partial Fixed-Point Logic* FO(S-PFP) and by showing that it has not more expressive power than FO(PFP), even though it will turn out to be a valuable tool in many considerations.

The formulas of FO(S-PFP) are obtained by replacing the PFP rule above by

$$\frac{\varphi_0, \dots, \varphi_m}{[\text{S-PFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}}$$

where $m \geq 0$, where for $i = 0, \dots, m$ the arity of X_i equals the length of \bar{x}_i , and where \bar{t} has the same length as \bar{x}_0 . As usual, the variables in each sequence \bar{x}_i are distinct, but the same variable may occur in distinct \bar{x}_i 's. A variable x is free in $[\text{S-PFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$, if x occurs in \bar{t} or if for at least one i , x is free in φ_i and not in \bar{x}_i . Introducing F^0, \dots, F^m by

$$F^i(X_0, \dots, X_m) := \{\bar{x}_i \mid \varphi_i(\bar{x}_i, X_0, \dots, X_m)\},$$

the meaning of $[\text{S-PFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t}$ is defined to be $\bar{t} \in F^0_{(\infty)}$.

Once more, for simplicity, we restrict ourselves to structures with at least two elements.

Theorem 8.3.4 $\text{FO}(\text{PFP}) \equiv \text{FO}(\text{S-PFP})$.

The proof of $\text{FO}(\text{S-PFP}) \leq \text{FO}(\text{PFP})$ is by induction on formulas, the main step being taken care of by the following lemma.

Lemma 8.3.5 (Simultaneous Induction for PFP) *Let*

$$\varphi_0(\bar{x}_0, X_0, \dots, X_m), \dots, \varphi_m(\bar{x}_m, X_0, \dots, X_m)$$

be formulas of $\text{FO}(\text{PFP})$, where X_i is k_i -ary and the length of \bar{x}_i is k_i . Set $k := \max\{k_0, \dots, k_m\} + (m + 1)$ and define the $\text{FO}(\text{PFP})$ -formula $\chi_J(z_1, \dots, z_k, Z)$ by⁸

$$\begin{aligned} \chi_J(z_1, \dots, z_k, Z) := & \exists v \exists w (\neg v = w \wedge \\ & ((\varphi_0(z_1, \dots, z_{k_0}, Z_- \tilde{v} w^0, \dots, Z_- \tilde{v} w^m) \wedge \delta_0(z_{k_0+1}, \dots, z_k, v, w)) \\ & \vee (\varphi_1(z_1, \dots, z_{k_1}, Z_- \tilde{v} w^0, \dots, Z_- \tilde{v} w^m) \wedge \delta_1(z_{k_1+1}, \dots, z_k, v, w)) \\ & \vdots \\ & \vee (\varphi_m(z_1, \dots, z_{k_m}, Z_- \tilde{v} w^0, \dots, Z_- \tilde{v} w^m) \wedge \delta_m(z_{k_m+1}, \dots, z_k, v, w))))). \end{aligned}$$

Then, for any new variable u ,

$$\models_{\text{fin}} [\text{S-PFP}_{\bar{x}_0, X_0, \dots, \bar{x}_m, X_m} \varphi_0, \dots, \varphi_m] \bar{t} \leftrightarrow \exists (\forall) u [\text{PFP}_{\bar{z}, Z} \chi_J] \bar{t} \bar{u}.$$

If $\varphi_0, \dots, \varphi_m$ are first-order formulas then χ_J is first-order. Moreover, given a structure and an assignment to variables, the simultaneous fixed-point of the PFP operation of the formula on the right side in the equivalence above exists iff the fixed-point of the PFP operation on the left side exists.

The lemma is a literal translation of the Lemma on Simultaneous Induction for LFP (cf. 8.2.3), and also the proof can be translated; of course, here we do not have the assumption that the formulas are positive in the second-order variables and hence, the same applies to χ_J . The additional claim on the existence of the fixed-point immediately follows from 8.2.1(a). \square

Also the proof of 8.2.4 can be literally translated and gives

⁸ For the definition of δ_i ; see page 179.

Lemma 8.3.6 *Every FO(PFP)-formula $[\text{PFP}_{\bar{y},Y} \varphi] \bar{t}$ with first-order φ is equivalent to a formula of the form*

$$\exists(\forall)u[\text{PFP}_{\bar{z},Z} \psi] \bar{u}$$

with first-order ψ . If φ is existential then ψ can be chosen existential, too. For a given structure and assignment, if F_∞^φ exists then so does F_∞^ψ . Furthermore, we may require that u does not occur free in ψ and that the free variables of $\exists(\forall)u[\text{PFP}_{\bar{z},Z} \psi] \bar{u}$ are among the free variables of $[\text{PFP}_{\bar{y},Y} \varphi] \bar{t}$. \square

In the proof of the Transitivity Lemma for LFP we have made essential use of the monotonicity of the operations in question, so it does not translate to the present case. However, it is easy to derive a version of the Transitivity Lemma for PFP (see 8.3.10 below), if we restrict ourselves to formulas whose fixed-points always exist. For this purpose call a formula of FO(PFP) *totally defined*, if for all its subformulas of the form $[\text{PFP}_{\bar{x},X} \psi] \bar{t}$ the fixed-point F_∞^ψ exists in every structure and for all assignments.

Exercise 8.3.7 Show, using Trahtenbrot’s theorem, that it is not decidable whether an FO(PFP)-formula is totally defined. \square

Nevertheless we can show:

Proposition 8.3.8 *For every FO(PFP)-formula φ there is an equivalent totally defined FO(PFP)-formula ψ . If φ has the form $[\text{PFP}_{\bar{x},X} \varphi'] \bar{t}$ with first-order φ' then ψ can be chosen of the form $[\text{PFP}_{\bar{x},X} \psi'] \bar{s}$ with first-order ψ' .*

For the proof we need the following

Lemma 8.3.9 *Assume M is finite. Let $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$. Then there is an $l \geq 1$ such that $F_l = F_{2 \cdot l}$.*

Proof. By finiteness of M , the sequence F_0, F_1, \dots must become periodic (cf. 7.1.1(a)). So there are n_0 and $l_0 \geq 1$ such that $F_m = F_{m+l_0}$ holds for all $m > n_0$. Choose s such that $l := s \cdot l_0 > n_0$. Then $F_l = F_{s \cdot l_0} = F_{s \cdot l_0 + s \cdot l_0} = F_{2 \cdot l}$. \square

Proof (of 8.3.8). The proof proceeds by induction on FO(PFP)-formulas, the main step being the PFP operator. So let $[\text{PFP}_{\bar{x},X} \chi] \bar{t}$ be given and suppose, by induction hypothesis, that χ is totally defined. Consider the function F^χ for a given structure \mathcal{A} and an assignment of values to the parameters. Set $X_n := F_n^\chi$ and $X_\infty := F_\infty^\chi$. Choose the smallest $l \geq 1$ such that $X_{2 \cdot l} = X_l$. We consider the sequences $(Y_n)_{n \geq 0}$, $(Z_n)_{n \geq 0}$, and $(W_n)_{n \geq 0}$, whose values are shown by the following table:

X	X_0	X_1	X_2	\dots	\dots	X_l	X_{l+1}	X_{l+2}	\dots
Y	X_0	X_1	X_2	\dots	\dots	X_l	X_∞	X_∞	\dots
Z	X_0	X_2	X_4	\dots	\dots	X_{2l}	X_∞	X_∞	\dots
W	\emptyset	A	A	\dots	\dots	A	A	A	\dots

In the first l steps (Y_n) proceeds as (X_n) and (Z_n) as ($X_{2,n}$), i.e., with twice the velocity of (X_n). When $n > 0$ (which is signaled by $W_n \neq \emptyset$) and $Y_n = Z_n$ then $Y_m = Z_m = X_\infty$ for $m > n$. So, all these new sequences have fixed-points (this guarantees that the formula obtained below is totally defined). The fixed-points are the simultaneous fixed-point of the operations given by the formulas φ_Y, φ_Z , and φ_W , where

$$\begin{aligned}\varphi_Y(\bar{y}, Y, Z, W) &:= ((W = \emptyset \vee Y \neq Z) \wedge \chi(\bar{y}, Y)) \vee (Y = \chi(-, Y) \wedge Y\bar{y})^9 \\ \varphi_Z(\bar{z}, Y, Z, W) &:= ((W = \emptyset \vee Y \neq Z) \wedge \chi(\bar{z}, \chi(-, Z))) \\ &\quad \vee (Z = \chi(-, Z) \wedge Z\bar{z}) \\ \varphi_W(w, Y, Z, W) &:= w = w.\end{aligned}$$

By the first two lines of the table we know that

$$\models_{\text{fin}} [\text{PFP}_{\bar{x}, X} \chi] \bar{t} \leftrightarrow [\text{S-PFP}_{\bar{y}, Y, \bar{z}, Z, w, W} \varphi_Y, \varphi_Z, \varphi_W] \bar{t}$$

and therefore, by the Lemma on Simultaneous Induction for PFP, we have for the corresponding χ_J that

$$\models_{\text{fin}} [\text{PFP}_{\bar{x}, X} \chi] \bar{t} \leftrightarrow [\text{PFP}_{\bar{v}, V} \chi_J] \bar{t}\bar{u}$$

and that the formula on the right side is totally defined. \square

Next we show an FO(PFP)-version of the Transitivity Lemma. It will be used in the proof of 8.3.3.

Lemma 8.3.10 (Transitivity for PFP) *Let $\varphi(\bar{x}, X, Y)$ and $\psi(\bar{y}, X, Y)$ be first-order and let $[\text{PFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)] \bar{y}$ be totally defined. (Moreover, assume that no individual variable free in $[\text{PFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)]$ gets into the scope of a corresponding quantifier or PFP operator in $(*)$.) Then the formula*

$$(*) \quad [\text{PFP}_{\bar{x}, X} \varphi(\bar{x}, X, [\text{PFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)])] \bar{t}$$

is equivalent to a totally defined formula

$$\exists(\forall)u[\text{PFP}_{\bar{z}, Z} \chi(\bar{z}, Z)] \bar{u}$$

with first-order χ .

Proof. Set $\rho := [\text{PFP}_{\bar{x}, X} \varphi(\bar{x}, X, [\text{PFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)])] \bar{t}$. We want to write the nested fixed-point as a simultaneous fixed-point. For this purpose let

$$\varphi'(\bar{x}, X, Y) := (Y \neq \psi(-, X, Y) \wedge X\bar{x}) \vee (Y = \psi(-, X, Y) \wedge \varphi(\bar{x}, X, Y))$$

(“if Y is not a fixed-point then X does not change, otherwise X goes one step further”),

⁹ $Y = \chi(-, Y)$ is an abbreviation for $\forall \bar{x}(Y\bar{x} \leftrightarrow \chi(\bar{x}, Y))$.

$$\psi'(\bar{y}, X, Y) := (Y \neq \psi(-, X, Y) \wedge \psi(\bar{y}, X, Y)) \vee (Y = \psi(-, X, Y) \wedge (X = \varphi(-, X, Y) \wedge Y\bar{y}))$$

(“if Y is not a fixed-point then Y goes one step further; otherwise Y is set back to \emptyset except X happens to be a fixed-point; then Y does not change”).

Since $[\text{PFP}_{\bar{y}, Y} \psi(\bar{y}, X, Y)]\bar{t}$ is totally defined, we immediately have

$$\models_{\text{fin}} \rho \leftrightarrow [\text{S-PFP}_{\bar{x}, X, \bar{y}, Y} \varphi', \psi']\bar{t}.$$

For the corresponding χ_J , which is first-order (cf. 8.3.5), we therefore get

$$\models_{\text{fin}} \rho \leftrightarrow [\text{PFP}_{\bar{z}, Z} \chi_J]\bar{t}\bar{u}.$$

To get a totally defined formula of the desired form, we first apply 8.3.8 to $[\text{PFP}_{\bar{z}, Z} \chi_J]\bar{t}\bar{u}$ and then 8.3.6. \square

Exercise 8.3.11 Let $\varphi(x, X, Y) := (x = c \vee Yx)$ and $\psi(y, X, Y) := \neg Yy$. Show that $[\text{PFP}_{x, X} \varphi(x, X, [\text{PFP}_{y, Y} \psi(y, X, Y)])]x$ is equivalent to the formula $x = c$, while the simultaneous fixed-point of φ' and ψ' defined according to the previous proof does not exist. Note that $[\text{PFP}_{y, Y} \psi(y, X, Y)]y$ is not totally defined. \square

We now get Theorem 8.3.3 in the following sharper form.

Theorem 8.3.12 *For every FO(PFP)-formula φ there is a first-order formula ψ such that*

$$\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)u[\text{PFP}_{\bar{z}, Z} \psi]\bar{u}$$

and $\exists(\forall)u[\text{PFP}_{\bar{z}, Z} \psi]\bar{u}$ is totally defined.

Proof. We proceed by induction on φ . The atomic case and the cases for \forall , \exists , and PFP are handled as in 8.2.7, now using the preceding lemma instead of 8.2.6. Concerning the negation step, suppose $\varphi = \neg\varphi_1$ and, by induction hypothesis, assume that

$$\varphi_1 = [\text{PFP}_{\bar{v}, V} \psi_1(\bar{v}, V)]\bar{t}$$

is totally defined and ψ_1 is first-order. Set $\rho(\bar{x}, X, Z) := \neg Z\bar{t}$ for dummy variables \bar{x} and X . Then $\varphi (= \neg[\text{PFP}_{\bar{v}, V} \psi_1(\bar{v}, V)]\bar{t})$ is equivalent to

$$[\text{PFP}_{\bar{x}, X} \rho(\bar{x}, X, [\text{PFP}_{\bar{v}, V} \psi_1(\bar{v}, V)])]\bar{t}$$

which, by the preceding lemma, is equivalent to a formula of the desired form. \square

We close this section by showing that in the normal form given in the preceding theorem the formula ψ on the right hand side can be chosen existential, a result we shall also obtain by different methods in Chapter 9. The following exercise gives the idea underlying the procedure to eliminate universal quantifiers.

Exercise 8.3.13 Show that in models of $\neg c = d$ the formulas $\forall xRx$ and

$$[\text{PFP}_{x,X} (x = c \vee (x = d \wedge \exists z\neg Rz)) \wedge ((Xc \wedge Xd) \rightarrow \neg x = x)]c$$

are equivalent. \square

Theorem 8.3.14 For every FO(PFP)-formula φ there is an existential first-order formula ρ such that $\exists(\forall)u[\text{PFP}_{\bar{y},Y} \rho] \tilde{u}$ is totally defined and

$$\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)u[\text{PFP}_{\bar{y},Y} \rho] \tilde{u}.$$

Proof. By the preceding theorem we can assume that $\varphi = [\text{PFP}_{\bar{y},Y} \psi] \tilde{u}$,¹⁰ where $[\text{PFP}_{\bar{y},Y} \psi] \tilde{u}$ is totally defined and where ψ is first-order and in prenex normal form. It suffices to show that if ψ is a Σ_{k+1} -formula (compare 1.B for the definition of Σ_n -formulas) with $k > 0$ then ψ can be replaced by a Σ_k -formula (and hence, by induction, by a Σ_1 -formula; the result then follows from 8.3.6). So assume that $\psi(\bar{y}, Y) = \exists \bar{z} \forall \bar{u} \psi'$ with $\psi' = \psi'(\bar{y}, \bar{z}, \bar{u}, Y) \in \Sigma_{k-1}$, and let $Y_n := F_n^\psi$. Define the sequences $(U_{(n)})$, $(V_{(n)})$, $(W_{(n)})$ simultaneously, using the formulas χ_U, χ_V, χ_W , respectively, where

$$\begin{aligned} \chi_U(\bar{y}, U, V, W) &:= W \neq \emptyset \wedge \exists \bar{z} \neg V \bar{y} \bar{z} \\ \chi_V(\bar{y} \bar{z}, U, V, W) &:= \neg \forall \bar{u} \psi'(\bar{y}, \bar{z}, \bar{u}, U) \\ \chi_W(x, U, V, W) &:= x = x. \end{aligned}$$

Note that $W_{(0)} = \emptyset$ and $W_{(n)} = \text{universe}$ for $n > 0$. Hence, we get $U_{(1)} = \emptyset$ and for $n \geq 2$:

$$\begin{aligned} U_{(n)} &= \{\bar{y} \mid \exists \bar{z} \neg V_{(n-1)} \bar{y} \bar{z}\} \\ &= \{\bar{y} \mid \exists \bar{z} \forall \bar{u} \psi'(\bar{y}, \bar{z}, \bar{u}, U_{(n-2)})\}. \end{aligned}$$

Since $Y_n = \{\bar{y} \mid \exists \bar{z} \forall \bar{u} \psi'(\bar{y}, \bar{z}, \bar{u}, Y_{n-1})\}$, a simple induction shows that $U_{(2n)} = Y_n$ and $U_{(2n+1)} = Y_n$; in particular, $(U_{(n)})$ eventually becomes constant. As $V_{(n)} = \{\bar{y} \bar{z} \mid \neg \forall \bar{u} \psi'(\bar{y}, \bar{z}, \bar{u}, U_{(n-1)})\}$, the sequence $(V_{(n)})$ becomes constant, too. Thus, the simultaneous fixed-point $(U_{(\infty)}, V_{(\infty)}, W_{(\infty)})$ exists and $U_{(\infty)} = Y_\infty$. Hence φ , that is, $[\text{PFP}_{\bar{y},Y} \psi] \tilde{u}$, is equivalent to $[\text{S-PFP}_{\bar{y},U,\bar{y}\bar{z},V,x,W} \chi_U, \chi_V, \chi_W] \tilde{u}$. Since χ_U, χ_V, χ_W are equivalent to a $\Sigma_1, \Sigma_k, \Sigma_1$ -formula, respectively, a closer look at the corresponding formula χ_J in 8.3.5 shows that it is equivalent to a Σ_k -formula and that, for a new variable v , we have

$$\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)v[\text{PFP}_{\bar{z},Z} \chi_J] \tilde{u}\tilde{v},$$

and hence, $\models_{\text{fin}} \varphi \leftrightarrow \exists(\forall)u[\text{PFP}_{\bar{z},Z} \chi_J] \tilde{u}\tilde{u}$. \square

Exercise 8.3.15 (A version of the game of life) Let $\mathcal{G} = (G, E^G, C^G)$ where (G, E^G) is a graph and $C^G \subseteq G$. Interpret C^G as the set of vertices hosting a live cell. Set $C_0 = \emptyset, C_1 = C^G$, and for $n > 1$, define C_n , the “ n -th generation”, by the following “reproduction rule”: For $a \in G$, $a \in C_n$ iff (i) or (ii) holds, where

¹⁰ Recall that a formula $\exists(\forall)x\chi(x)$ is equivalent to $\chi(t)$ for any term t .

- (i) $a \in C_{n-1}$ and, in C_{n-1} , a has exactly two neighbours hosting a live cell (that is, there are exactly two $b \in C_{n-1}$ such that $E^G ab$);
- (ii) in C_{n-1} , a has exactly three neighbours hosting a live cell.

Show that there is a sentence φ of FO(PFP) such that φ holds in a structure \mathcal{G} as above iff for some $n \geq 1$, $C_n = \emptyset$ (the species dies out). Give an example of a structure \mathcal{G} such that $(C_n)_{n \geq 0}$ has no fixed-point. \square

Exercise 8.3.16 (a) Assume that M is finite and $F : \text{Pow}(M) \rightarrow \text{Pow}(M)$. Define

$$F_{\text{true}} := \{a \in M \mid \exists n_0 \forall m \geq n_0 : a \in F_m\},$$

$$F_{\text{false}} := \{a \in M \mid \exists n_0 \forall m \geq n_0 : a \notin F_m\},$$

and $F_{\text{und}} := M \setminus (F_{\text{true}} \cup F_{\text{false}})$, the set of *undefined* points. Introduce the class of formulas of the logic FO(3-PFP) as that of FO(PFP) at the beginning of this section replacing, however, the last rule of that calculus by

$$\frac{\varphi}{[t\text{-PFP}_{\bar{x}, X} \varphi] \bar{t}}, \quad \frac{\varphi}{[f\text{-PFP}_{\bar{x}, X} \varphi] \bar{t}}, \quad \frac{\varphi}{[u\text{-PFP}_{\bar{x}, X} \varphi] \bar{t}},$$

the meaning of these formulas being

$$“\bar{t} \in F_{\text{true}}^\varphi”, \quad “\bar{t} \in F_{\text{false}}^\varphi”, \quad \text{and} \quad “\bar{t} \in F_{\text{und}}^\varphi”,$$

respectively. Show that $\text{FO}(3\text{-PFP}) \equiv \text{FO}(\text{PFP})$. Hint: Argue as in the proof of 8.3.8, using, e.g., for the truth set, an additional sequence $(V_n)_{n \geq 0}$ given by $V_0 = X_0, V_1 = X_1, \dots, V_l = X_l, V_{l+1} = V_l \cap X_{l+1}, V_{l+2} = V_{l+1} \cap \bar{X}_{l+2}, \dots$

(b) For the game associated with digraphs and $\varphi(x, X) := \exists y(-Xy \wedge Exy)$ the formulas $[t\text{-PFP}_{x, X} \varphi] x$, $[f\text{-PFP}_{x, X} \varphi] x$, and $[u\text{-PFP}_{x, X} \varphi] x$ express “ x is won for I”, “ x is won for II”, and “ x is drawn”, respectively.

Notes 8.3.17 Partial fixed-point logic goes back to [2]. The presentation of this section owes much to Grohe [61].

8.4 Fixed-Point Logics and $L_{\infty\omega}^\omega$

In the first part of this section we show that the fixed-point logics introduced so far are contained in $L_{\infty\omega}^\omega$. Therefore, when exploring the expressive power of fixed-point logics, we can prove nonaxiomatizability results by using nonaxiomatizability results for $L_{\infty\omega}^\omega$. The second part is concerned with the question whether $\text{FO}(\text{LFP}) \equiv \text{FO}(\text{PFP})$. We shall see that, when studying this problem, we can restrict ourselves to ordered structures. Thus, by the results of Chapter 7, the problem $\text{FO}(\text{LFP}) \equiv \text{FO}(\text{PFP})$ is equivalent to the question whether the complexity classes PTIME and PSPACE coincide.

Throughout this section, for simplicity, we only consider relational vocabularies. The following lemma shows a way how to define the stages of a fixed-point construction in first-order logic.

Proposition 8.4.1 *Let $\varphi(\bar{x}, X)$ be a first-order formula where all variables are among v_1, \dots, v_k and X . Suppose X is s -ary and $\bar{x} = x_1 \dots x_s$ (with x_1, \dots, x_s among v_1, \dots, v_k). Then, for every n , there is a formula $\varphi^n(\bar{x})$ in FO^{k+s} defining the stage F_n^φ .*

Proof. Let $y_1 = v_{k+1}, \dots, y_s = v_{k+s}$. Then $\varphi^n(\bar{x})$ can be defined inductively by

$$\begin{aligned}\varphi^0(\bar{x}) &:= \neg x_1 = x_1, \\ \varphi^{n+1}(\bar{x}) &:= \varphi(\bar{x}, X) \frac{\exists \bar{y}(\bar{y} = _ \wedge \exists \bar{x}(\bar{x} = \bar{y} \wedge \varphi^n(\bar{x})))}{X_}\end{aligned}$$

(i.e., we replace in $\varphi(\bar{x}, X)$ each occurrence of an atomic subformula of the form $X\bar{t}$ by $\exists \bar{y}(\bar{y} = \bar{t} \wedge \exists \bar{x}(\bar{x} = \bar{y} \wedge \varphi^n(\bar{x})))$; note that some variables of \bar{x} may occur in \bar{t}). \square

As a corollary we get:

Theorem 8.4.2 $\text{FO}(\text{PFP}) \leq L_{\infty\omega}^\omega$.

Proof. By 8.3.12, it suffices to show for first-order φ that $[\text{PFP}_{\bar{x}, X} \varphi] \bar{t}$ is equivalent to a formula of $L_{\infty\omega}^\omega$.

So suppose that $\varphi, k, s, X, \bar{x}, \bar{y}$ are as in the preceding proposition and that the variables in \bar{t} are among v_1, \dots, v_k . Then $[\text{PFP}_{\bar{x}, X} \varphi] \bar{t}$ is equivalent to the $L_{\infty\omega}^{k+s}$ -formula

$$\bigvee_{n \geq 0} (\forall \bar{x}(\varphi^n(\bar{x}) \leftrightarrow \varphi^{n+1}(\bar{x})) \wedge \varphi^n(\bar{t}))$$

where, to stay within $L_{\infty\omega}^{k+s}$, we take $\exists \bar{y}(\bar{y} = \bar{t} \wedge \exists \bar{x}(\bar{x} = \bar{y} \wedge \varphi^n(\bar{x})))$ for $\varphi^n(\bar{t})$. \square

Exercise 8.4.3 Show that $\forall x \forall y [\text{LFP}_{y, Y}(y = x \vee \exists x(Yx \wedge Exy))]y$ is not equivalent to a formula of $L_{\infty\omega}^2$. (Hint: Use part (b) of 3.3.6). \square

The preceding theorem allows us to translate some results from $L_{\infty\omega}^\omega$ to $\text{FO}(\text{PFP})$. We give two examples. By 3.3.25(b) we get

Proposition 8.4.4 *For a relational vocabulary containing only unary relation symbols the expressive power of $\text{FO}(\text{PFP})$ and FO coincide; in particular, the class EVEN of sets of even cardinality is not axiomatizable in $\text{FO}(\text{PFP})$.* \square

The last statement of the proposition shows that neither $\text{FO}(\text{LFP})$ strongly captures PTIME nor $\text{FO}(\text{PFP})$ strongly captures PSPACE (cf. page 157). From the 0–1 laws for $L_{\infty\omega}^\omega$ obtained in Chapter 4 we get for example:

Theorem 8.4.5 FO(PFP) *satisfies the labeled and the unlabeled 0–1 law.* \square

Now we turn to the question whether $\text{FO(LFP)} \equiv \text{FO(PFP)}$. We first recall some definitions: In subsection 3.3.2, to every τ -structure \mathcal{A} and to every $s \geq 1$ we have assigned a τ/s -structure \mathcal{A}/s which captures the $L_{\infty\omega}^s$ -theory of \mathcal{A} . Recall that \mathcal{A}/s consists of the classes of the equivalence relation \sim on A^s given by

$$\bar{a} \sim \bar{b} \quad \text{iff} \quad \bar{a} \text{ and } \bar{b} \text{ satisfy the same } L_{\infty\omega}^s\text{-formulas in } \mathcal{A}$$

and recall that the vocabulary τ/s and the structures \mathcal{A}/s are given as follows:

- for every k -ary $R \in \tau \cup \{=\}$ and any i_1, \dots, i_k with $1 \leq i_1, \dots, i_k \leq s$ the vocabulary τ/s contains a unary relation symbol $R_{i_1 \dots i_k}$ with interpretation

$$R_{i_1 \dots i_k}^{A/s} := \{[\bar{a}] \mid \bar{a} \in A^s, R^{\mathcal{A}} a_{i_1} \dots a_{i_k}\};$$

- for $i = 1, \dots, s$ the vocabulary τ/s contains a binary relation symbol S_i with

$$S_i^{A/s} := \{([\bar{a}], [\bar{a}']) \mid \bar{a}, \bar{a}' \in A^s, \text{there is } a \in A \text{ such that } [\bar{a}'] = [\bar{a} \frac{a}{i}]\}.$$

Since $\text{FO(PFP)} \leq L_{\infty\omega}^\omega$, for any $\varphi \in \text{FO(PFP)}[\tau]$ we can choose an s such that $\varphi = \varphi(v_1, \dots, v_s)$ is equivalent to an $L_{\infty\omega}^s$ -formula, thus

$$\mathcal{A} \models \varphi[\bar{a}] \text{ and } \bar{a} \sim \bar{b} \text{ imply } \mathcal{A} \models \varphi[\bar{b}].$$

Therefore, any FO(PFP) definable relation on \mathcal{A} is the union of some equivalence classes. This will enable us to show that the FO(LFP) (and the FO(PFP)) definable relations on \mathcal{A} and on the \mathcal{A}/s correspond to each other. Since \mathcal{A}/s has an FO(LFP) definable ordering, we thus shall be able to translate the question whether $\text{FO(LFP)} \equiv \text{FO(PFP)}$ from arbitrary structures to ordered ones.

Fix s . We start by showing that \sim is definable in FO(LFP) (cf. 8.1.8). Let \bar{a} and \bar{b} range over A^s . For $j \geq 0$ define \sim_j on A^s by induction:

$$\begin{aligned} \bar{a} \sim_0 \bar{b} & \quad \text{iff } \bar{a} \text{ and } \bar{b} \text{ satisfy the same atomic formulas in } \mathcal{A} \\ \bar{a} \sim_{j+1} \bar{b} & \quad \text{iff } \bar{a} \sim_0 \bar{b} \text{ and for all } i = 1, \dots, s \text{ and all } a \in A \text{ (} b \in A \text{)} \\ & \quad \text{there is } b \in A \text{ (} a \in A \text{) such that } \bar{a} \frac{a}{i} \sim_j \bar{b} \frac{b}{i} \end{aligned}$$

(in the terminology of subsection 3.3.1 we have $\bar{a} \sim_j \bar{b}$ iff the duplicator has a winning strategy in the pebble game $G_j^s(\mathcal{A}, \bar{a}, \mathcal{A}, \bar{b})$ with s pebbles and j moves). Clearly, $\sim_0 \supseteq \sim_1 \supseteq \dots$, so that $\sim_l = \sim_{l+1}$ for some l .¹¹ For such an l we have $\sim_l = \sim$. For the complements $\not\sim_j$ of \sim_j we have $\not\sim_0 \subseteq \not\sim_1 \subseteq \dots$

¹¹ In subsection 3.3.3 the least such l was called the s -rank $r(\mathcal{A})$.

They are the stages $F_1^{\varphi_{\neq}}, F_2^{\varphi_{\neq}}, \dots$, where $\varphi_{\neq}(x_1, \dots, x_s, y_1, \dots, y_s, Z)$ is the following formula positive in Z :

$$\bigvee_{\substack{\psi \in \text{FO}^s \\ \psi \text{ atomic}}} (\psi(\bar{x}) \leftrightarrow \neg\psi(\bar{y})) \vee \bigvee_{1 \leq i \leq s} (\exists x_i \forall y_i Z \bar{x} \bar{y} \vee \exists y_i \forall x_i Z \bar{x} \bar{y}).$$

Altogether we get

Lemma 8.4.6 *There is a τ -formula $\varphi_{\sim}(\bar{x}, \bar{y})$ in $\text{FO}(\text{LFP})$ with variables among $x_1, \dots, x_s, y_1, \dots, y_s$ such that for all τ -structures \mathcal{A} and $\bar{a}, \bar{b} \in A^s$,*

$$\bar{a} \sim \bar{b} \quad \text{iff} \quad \mathcal{A} \models \varphi_{\sim}[\bar{a}, \bar{b}].$$

Proof. Take as $\varphi_{\sim}(\bar{x}, \bar{y})$ the formula $\neg[\text{LFP}_{\bar{x}\bar{y}, Z} \varphi_{\neq}(\bar{x}, \bar{y}, Z)] \bar{x} \bar{y}$. \square

By this lemma there is an $\text{FO}(\text{LFP})$ definition of \mathcal{A}/s in \mathcal{A} and hence, any $\text{FO}(\text{LFP})$ statement on \mathcal{A}/s can be translated to an $\text{FO}(\text{LFP})$ statement on \mathcal{A} .

Proposition 8.4.7 *For any $\varphi(x_1, \dots, x_k)$ in $\text{FO}(\text{LFP})[\tau/s]$ whose free variables are among x_1, \dots, x_k there is a formula $\tilde{\varphi}(x_{11}, \dots, x_{1s}, \dots, x_{k1}, \dots, x_{ks})$ in $\text{FO}(\text{LFP})[\tau]$ such that for all τ -structures \mathcal{A} and $\bar{a}_1, \dots, \bar{a}_k$ in A^s ,*

$$\mathcal{A}/s \models \varphi[[\bar{a}_1], \dots, [\bar{a}_k]] \quad \text{iff} \quad \mathcal{A} \models \tilde{\varphi}[\bar{a}_1, \dots, \bar{a}_k].$$

The same is true if $\text{FO}(\text{LFP})$ is replaced by $\text{FO}(\text{PFP})$ everywhere.

Proof. For a variable y let \bar{y} be $y_1 \dots y_s$. The definition of $\tilde{\varphi}$ is by induction on φ . We give the main steps:

$$\begin{array}{l} \frac{\varphi}{y = z} \\ R_{i_1 \dots i_l} y \\ S_i y z \\ X y_1 \dots y_r \\ \exists y \psi \\ [\text{LFP}_{y_1 \dots y_r, X} \psi] z_1 \dots z_r \end{array} \qquad \begin{array}{l} \frac{\tilde{\varphi}}{\varphi_{\sim}(\bar{y}, \bar{z})} \\ R y_{i_1} \dots y_{i_l} \\ \exists y_i \varphi_{\sim}(\bar{y}, \bar{z}) \\ X' \bar{y}_1 \dots \bar{y}_r \\ \exists y_1 \dots \exists y_s \tilde{\psi} \\ [\text{LFP}_{\bar{y}_1 \dots \bar{y}_r, X'} \tilde{\psi}] \bar{z}_1 \dots \bar{z}_r \end{array}$$

where X' has s -times the arity of X . \square

We come to a translation in the opposite direction:

Proposition 8.4.8 *For any $\text{FO}(\text{LFP})[\tau]$ -formula φ without free second-order variables there is an $s \geq 1$ and an $\text{FO}(\text{LFP})[\tau/s]$ -formula $\varphi^*(v)$ with at most the free variable v such that the individual variables in φ are among v_1, \dots, v_s and such that for all τ -structures \mathcal{A} and $\bar{a} \in A^s$,*

$$\mathcal{A} \models \varphi[a_1, \dots, a_s] \quad \text{iff} \quad \mathcal{A}/s \models \varphi^*[[\bar{a}]].$$

If φ is a sentence then φ^ can be chosen to be a sentence.*

The same is true for $\text{FO}(\text{PFP})$ instead of $\text{FO}(\text{LFP})$.

Proof. By Theorem 8.1.14, we may assume that the LFP operator occurs only once in φ , w.l.o.g. in the form $[\text{LFP}_{v_1 \dots v_r, X} \psi] \bar{t}$. Moreover, we may assume that the occurrences of X have the form $Xv_1 \dots v_r$ (if $Xz_1 \dots z_r$ occurs, replace it by $\exists \bar{y}(\bar{y} = \bar{z} \wedge \exists \bar{v}(\bar{v} = \bar{y} \wedge X\bar{v}))$ where \bar{y} are new variables and $\bar{v} = v_1 \dots v_r$). Treating \bar{t} similarly, we may assume that the LFP operator in φ has the form $[\text{LFP}_{v_1 \dots v_r, X} \psi] v_1 \dots v_r$. Let the variables in φ be among v_1, \dots, v_s . Finally, we may assume that $r = s$; otherwise, we replace X by an s -ary variable Y and $[\text{LFP}_{v_1 \dots v_r, X} \psi(X)] v_1 \dots v_r$ by $[\text{LFP}_{v_1 \dots v_s, Y} \psi(Y_{-v_{r+1} \dots v_s})] v_1 \dots v_s$.

The definition of φ^* is by induction on φ , the main steps being (we use a unary relation variable V):

$$\frac{\varphi}{v_i = v_j} \qquad \frac{\varphi^*}{=_{ij} v}$$

$$\frac{}{Rv_{i_1} \dots v_{i_k}} \qquad \frac{}{R_{i_1 \dots i_k} v}$$

$$\frac{}{Xv_1 \dots v_s} \qquad \frac{}{Vv}$$

$$\frac{}{\exists v_i \psi} \qquad \frac{}{\exists w (S_i v w \wedge \psi^*(w))}$$

$$\frac{}{[\text{LFP}_{v_1 \dots v_s, X} \psi] v_1 \dots v_s} \qquad \frac{}{[\text{LFP}_{v, V} \psi^*] v}$$

For a sentence φ pass to the sentence $\exists v \varphi^*$. □

The next proposition shows that from the point of view of FO(LFP) the structures \mathcal{A}/s can be regarded as ordered structures.

Proposition 8.4.9 *There is an FO(LFP)[τ/s]-formula $\psi_{<}(u, v)$ such that for all τ -structures \mathcal{A} ,*

$$\psi_{<}^{\mathcal{A}/s} := \{([\bar{a}], [\bar{b}]) \mid \mathcal{A}/s \models \psi_{<}([\bar{a}], [\bar{b}])\}$$

is an ordering on \mathcal{A}/s .

Proof. Let \mathcal{A} be a τ -structure and let \sim_j be the equivalence relations introduced before 8.4.6. Assume $\sim_l = \sim$. For $j \geq 0$ we define partial orderings (that means, irreflexive and transitive relations) $<_j$ on \mathcal{A}/s such that

$$(1) \qquad <_0 \subseteq <_1 \subseteq <_2 \dots$$

$$(2) \qquad ([\bar{a}] <_j [\bar{b}] \text{ or } [\bar{b}] <_j [\bar{a}]) \quad \text{iff} \quad \bar{a} \not\sim_j \bar{b}$$

(thus, $<_j$ induces an ordering on the equivalence classes of \sim_j). Then $<_l$ will be the intended ordering.

Let χ_1, \dots, χ_m be an enumeration – say in lexicographic order – of the atomic τ -formulas in the variables v_1, \dots, v_s and let χ_i^* be the τ/s -formula associated to χ_i in Proposition 8.4.8. As $<_0$ we take $<_0 := \psi_0^{\mathcal{A}/s}(-, -)$ where

$$\psi_0(x, y) := \bigvee_{1 \leq i \leq m} (\chi_i^*(x) \wedge \neg \chi_i^*(y) \wedge \bigwedge_{1 \leq j < i} (\chi_j^*(x) \leftrightarrow \chi_j^*(y))).$$

Clearly, $<_0$ is a partial ordering that satisfies (2) for $j = 0$.

Now suppose that $<_j$ has already been defined. To extend $<_j$ to $<_{j+1}$, assume that $\bar{a} \sim_j \bar{b}$ and $\bar{a} \not\sim_{j+1} \bar{b}$. For $1 \leq i \leq s$ set $M(\bar{a}, i) := \{[\bar{a}_i^a]_j \mid a \in A\}$ and $M(\bar{b}, i) := \{[\bar{b}_i^b]_j \mid b \in A\}$, where $[\bar{c}]_j$ denotes the equivalence class of \bar{c} modulo \sim_j . Since $\bar{a} \not\sim_{j+1} \bar{b}$, there is a smallest number i such that $M(\bar{a}, i) \neq M(\bar{b}, i)$. Set $[\bar{a}] <_{j+1} [\bar{b}]$ if the $<_j$ -minimal element of $(M(\bar{a}, i) \setminus M(\bar{b}, i)) \cup (M(\bar{b}, i) \setminus M(\bar{a}, i))$ belongs to $M(\bar{a}, i)$; otherwise, set $[\bar{b}] <_{j+1} [\bar{a}]$. One easily verifies that $<_{j+1}$ is a partial ordering (extending $<_j$) and that (2) holds for $j + 1$.

Clearly, $< := <_l$ is a total ordering on A/s . It is the fixed-point of the increasing relations $<_j$ and is thus definable in \mathcal{A}/s by $[\text{IFP}_{uv, X} \chi] uv$ where, with $x \equiv y$ for $\neg Xxy \wedge \neg Xyx$,

$$\begin{aligned} \chi(u, v, X) := & \psi_0(u, v) \vee \bigvee_{1 \leq i \leq s} (\exists u' (S_i u u' \wedge \forall v' (S_i v v' \rightarrow \\ & (X u' v' \vee (X v' u' \wedge \exists u'' (S_i u u'' \wedge u'' \equiv v'))))) \\ & \wedge \bigwedge_{1 \leq k < i} (\forall u' (S_k u u' \rightarrow \exists v' (S_k v v' \wedge u' \equiv v')) \\ & \wedge \forall v' (S_k v v' \rightarrow \exists u' (S_k u u' \wedge u' \equiv v'))). \end{aligned}$$

□

We are now able to show the main result of this section.

Theorem 8.4.10 *The following are equivalent:*

- (i) $\text{FO}(\text{PFP}) \equiv \text{FO}(\text{LFP})$.
- (ii) $\text{FO}(\text{PFP}) \equiv \text{FO}(\text{LFP})$ on ordered structures.

Proof. The direction from (i) to (ii) is clear. Assume $\text{FO}(\text{PFP}) \equiv \text{FO}(\text{LFP})$ on ordered structures. Let τ be a vocabulary and suppose φ is an $\text{FO}(\text{PFP})[\tau]$ -sentence. We must show that there is an $\text{FO}(\text{LFP})[\tau]$ -sentence with the same finite models. Choose $s \geq 1$ and $\varphi^* \in \text{FO}(\text{PFP})[\tau/s]$ according to 8.4.8. By (ii) there is an $\text{FO}(\text{LFP})[\tau/s \dot{\cup} \{<\}]$ -sentence ψ such that

$$\text{“} < \text{ is an ordering”} \models \varphi^* \leftrightarrow \psi.$$

Let $\chi := \psi \frac{\psi_{<(\dots-)}}{\dots < -}$, i.e., χ is obtained from ψ by replacing any atomic part of the form $x < y$ by $\psi_{<}(x, y)$, where $\psi_{<}(u, v)$ is the formula of the preceding proposition. Hence, $\chi \in \text{FO}(\text{LFP})[\tau/s]$. Choose $\tilde{\chi} \in \text{FO}(\text{LFP})[\tau]$ according to 8.4.7. Then, for any τ -structure \mathcal{A} we have

$$\begin{aligned} \mathcal{A} \models \varphi & \quad \text{iff} & \quad \mathcal{A}/s \models \varphi^* \\ & \quad \text{iff} & \quad (\mathcal{A}/s, \psi_{<}^{\mathcal{A}/s}) \models \psi \\ & \quad \text{iff} & \quad \mathcal{A}/s \models \chi \\ & \quad \text{iff} & \quad \mathcal{A} \models \tilde{\chi}. \end{aligned}$$

□

With 7.5.7 we now get

Corollary 8.4.11 $\text{FO(LFP)} \equiv \text{FO(PFP)}$ iff $\text{PTIME} = \text{PSPACE}$. \square

The ordered structure $(\mathcal{A}/s, <)$ is isomorphic to exactly one structure of the form $(\{0, \dots, n\}, \dots, <)$, the *numerical s -invariant* of \mathcal{A} . By the results above, a query on the class of τ -structures is definable in FO(LFP) iff for all sufficiently large s , the corresponding query on the numerical s -invariants is definable in FO(LFP) . Therefore, by 7.5.2, a query on the class of τ -structures is definable in FO(LFP) iff the corresponding query on the numerical invariants is in PTIME . That is, on arbitrary finite structures exactly those queries are definable in FO(LFP) which correspond to PTIME queries on the numerical invariants. The same applies to FO(PFP) and PSPACE .

Exercise 8.4.12 A τ -structure \mathcal{A} is *s -rigid* if $\tilde{a} \not\sim \tilde{b}$ for all $a, b \in A$, $a \neq b$ (where \sim is the equivalence relation leading to \mathcal{A}/s and \tilde{a}, \tilde{b} are the constant sequences of length s). A class K of structures is *s -rigid*, if every structure in K is s -rigid. Show:

- (a) \mathcal{A} is rigid iff \mathcal{A} is s -rigid for some s .
- (b) The following are equivalent:
 - (i) K is s -rigid for some s .
 - (ii) There is a formula in FO(LFP) defining an ordering on K (i.e., on every structure in K).
 - (iii) There is a formula in $L_{\infty\omega}^\omega$ defining an ordering on K .

Hint: (i) \Rightarrow (ii): If the τ -structure \mathcal{A} is s -rigid, then $\chi^A(-, -)$ with $\chi(x, y) = \tilde{\psi}_{<}(x, \dots, x, y, \dots, y)$ is an ordering, where $\tilde{\psi}_{<}(\bar{u}, \bar{v})$ is the formula associated with the τ/s -formula $\psi_{<}(u, v)$ of 8.4.9 according to 8.4.7. – (iii) \Rightarrow (i): If $\varphi(x, y) \in L_{\infty\omega}^\omega$ defines an ordering, use formulas as defined in part (a) of 3.3.1, here built up from $\varphi(x, y)$ instead of $x < y$, to show that K is s -rigid for some s .

Conclude: If K is an s -rigid class of structures, then the PTIME -computable queries on K are exactly the FO(LFP) definable ones.

In contrast to (a) there is a first-order axiomatizable class of rigid structures that is not s -rigid for any s (see [77]). \square

Exercise 8.4.13 Let K be a class of structures. We say that K is *fixed-point bounded*, if for any first-order formula $\varphi(\bar{x}, X)$ positive in X with free variables among \bar{x}, X , there is an m_0 such that

$$K \models \forall \bar{x} (\varphi^{m_0+1}(\bar{x}) \rightarrow \varphi^{m_0}(\bar{x})),$$

where $\varphi^m(\bar{x})$ is a first-order formula defining F_m^φ (see 8.4.1). Show the equivalence of (i) – (iii):

- (i) K is fixed-point bounded.
- (ii) K is bounded (in the sense of Theorem 3.3.24).
- (iii) On K , every $L_{\infty\omega}^\omega$ -formula is equivalent to an FO-formula.

Hint: For (i) \Rightarrow (ii) apply the hypothesis to the formula $\varphi_{\neq}(\bar{x}, \bar{y}, X)$ introduced before 8.4.6. The equivalence of (ii) and (iii) was shown in 3.3.24. For the direction from (iii) to (i) suppose that $\varphi(\bar{x}, X)$ is as above and witnesses that there is no m_0 as claimed. For $M \subseteq \mathbb{N}$ set $\varphi_M := \bigvee_{m \in M} (\varphi^{m+1}(\bar{x}) \wedge \neg \varphi^m(\bar{x}))$. Then, φ_M and φ_L are not equivalent on K for $L \neq M$.

In (iii), one cannot replace $L_{\infty\omega}^\omega$ by $\text{FO}(\text{LFP})$ (see [74]). \square

8.4.1 The Logic $\text{FO}(\text{PFP}_{\text{PTIME}})$

While least fixed-point operations reach their fixed-point in polynomially many steps, partial fixed-point operations may need exponentially many steps. The intermediate logic $\text{FO}(\text{PFP}_{\text{PTIME}})$ is the fragment of $\text{FO}(\text{PFP})$ consisting of those formulas for which every occurrence of the PFP operator reaches a fixed-point in polynomially many steps in all structures. Clearly,

$$\text{FO}(\text{LFP}) \leq \text{FO}(\text{PFP}_{\text{PTIME}}).$$

Exercise 8.4.16 shows that $\text{FO}(\text{LFP}) \equiv \text{FO}(\text{PFP}_{\text{PTIME}})$ for arbitrary structures is quite unlikely. On the other hand we have

Theorem 8.4.14 *On ordered structures, $\text{FO}(\text{LFP}) \equiv \text{FO}(\text{PFP}_{\text{PTIME}})$.*

Proof. The claim is clear by the results of Chapter 7: By definition, every query definable in $\text{FO}(\text{PFP}_{\text{PTIME}})$ can be evaluated in polynomial time and is thus definable in $\text{FO}(\text{LFP})$. Nevertheless, for later purposes we sketch a direct proof. Let $[\text{PFP}_{\bar{x}, X} \varphi(\bar{x}, X)] \bar{t}$ be given and suppose that, for some polynomial p , the sequence $(X_l)_{l \geq 0}$ (with $X_l := F_l^\varphi$) gets constant after at most p (cardinality of the universe) many steps. The idea of how to express this fixed-point in $\text{FO}(\text{LFP})$ is simple: We code the stages $(X_l)_{l \leq p(\dots)}$ by a relation variable Y with fixed-point Y_∞ where

$$Y_\infty m \bar{x} \quad \text{iff} \quad m \leq p(\dots) \wedge X_m \bar{x}.$$

The stages of Y_0, Y_1, Y_2, \dots will form the increasing sequence

$$\begin{aligned} & \{(0, \bar{x}) \mid \bar{x} \in X_0\} \quad (= \emptyset), \\ & \{(0, \bar{x}) \mid \bar{x} \in X_0\} \cup \{(1, \bar{x}) \mid \bar{x} \in X_1\}, \\ & \{(0, \bar{x}) \mid \bar{x} \in X_0\} \cup \{(1, \bar{x}) \mid \bar{x} \in X_1\} \cup \{(2, \bar{x}) \mid \bar{x} \in X_2\}, \dots \end{aligned}$$

The arithmetics needed for this encoding is definable in $\text{FO}(\text{LFP})$. In fact, by 8.1.11 there are $k \geq 1$ and formulas $S^k \bar{u} \bar{v}$ and $\bar{v} = p(x + 1)$ in $\text{FO}(\text{LFP})$ expressing (in the corresponding number representation) that \bar{v} is the successor of \bar{u} and that \bar{v} equals $p(x + 1)$, respectively. Therefore, as a formula equivalent to $[\text{PFP}_{\bar{x}, X} \varphi(\bar{x}, X)] \bar{t}$ we can take

$$\begin{aligned} \exists \bar{w} (\bar{w} = p(\max + 1) \wedge [\text{IFP}_{\bar{u} \bar{x}, Y} ((S^k \widetilde{\min} \bar{u} \wedge \varphi(\bar{x}, \emptyset)) \\ \vee \exists \bar{v} \bar{y} (Y \bar{v} \bar{y} \wedge \bar{v} < \bar{w} \wedge S^k \bar{v} \bar{u} \wedge \varphi(\bar{x}, Y \bar{v}_-))]) \bar{w} \bar{t}). \end{aligned}$$

\square

Exercise 8.4.15 ($\text{FO(LFP)} < \text{L}_{\infty\omega}^\omega \cap \text{PTIME}$) Let $\tau = \{E, U\}$ and let $\tau' = \{<, U\}$ with unary U and binary $<, E$. Denote by STG the class of (binary) “strings”, i.e., the class of *ordered* τ' -structures. A finite τ -structure $\mathcal{T} = (T, E^T, U^T)$ is a *labeled complete binary tree* if

- (T, E^T) is a tree, where every element has out-degree 0 or 2 and all leaves have the same depth (cf. 8.1.7).
- For any $a, b \in T$ of the same depth we have $U^A a$ iff $U^A b$.

Denote by LCT the class of labeled complete binary trees. Show:

- LCT is axiomatizable in FO(LFP) (use 8.1.7).
- If $\mathcal{T}_1, \mathcal{T}_2 \in \text{LCT}$ and $\mathcal{T}_1 \equiv_{\text{L}_{\infty\omega}^2} \mathcal{T}_2$ then $\mathcal{T}_1 \cong \mathcal{T}_2$.
- Any class K with $K \subseteq \text{LCT}$ is axiomatizable in $\text{L}_{\infty\omega}^\omega$.

There is a natural correspondence between strings and labeled complete binary trees obtained by encoding the same U -information with the i -th bit of the string as with the i -th level of the tree. More formally, if $\varphi_i(x) \in \text{FO}[\tau']$ expresses that x is the i -th element of $<$ and $\psi_i(x) \in \text{FO}[\tau]$ expresses that x has depth i , then for $\mathcal{A} \in \text{STG}$ and $\mathcal{T} \in \text{LCT}$ set $\mathcal{A} \approx \mathcal{T}$ iff for all $i \geq 0$,

- $\mathcal{A} \models \exists x \varphi_i(x)$ iff $\mathcal{T} \models \exists x \psi_i(x)$
- $\mathcal{A} \models \forall x (\varphi_i(x) \rightarrow Ux)$ iff $\mathcal{T} \models \forall x (\psi_i(x) \rightarrow Ux)$.

Show

- If $H \subseteq \text{LCT}$ is axiomatizable in FO(LFP) then $\text{STG}(H) := \{\mathcal{A} \mid \mathcal{A} \approx \mathcal{T} \text{ for some } \mathcal{T} \in H\}$ is axiomatizable in FO(LFP) . Hint: Show for $\mathcal{T} \in \text{LCT}$ and $\bar{a}, \bar{b} \in T^s$ that

$$(\mathcal{T}, \bar{a}) \cong (\mathcal{T}, \bar{b}) \quad \text{iff} \quad (d(a_i \wedge a_j))_{1 \leq i < j \leq s} = (d(b_i \wedge b_j))_{1 \leq i < j \leq s},$$

where $d(c)$ is the depth of c and where $a \wedge b$ denotes the element c of least depth with $c \leq a$ and $c \leq b$ in the induced partial ordering (cf. 8.1.7).

Thus s -tuples of elements of the tree can be encoded by $\frac{s(s+1)}{2}$ -tuples in the corresponding string. The set of the latter tuples being definable in FO, this can be used to translate the FO(LFP) -axiomatization of H into an FO(LFP) -axiomatization of $\text{STG}(H)$.

- If $K \subseteq \text{STG}$ is in $\text{TIME}(2^n)$ (that is, K is accepted by a 2^n time-bounded machine), then $\text{LCT}(K)_<$ is in PTIME , where $\text{LCT}(K) := \{\mathcal{T} \mid \mathcal{A} \approx \mathcal{T} \text{ for some } \mathcal{A} \in K\}$.

It is a well-known result of complexity theory that there are classes $K \subseteq \text{STG}$ that are in $\text{TIME}(2^n)$ but not in PTIME . Conclude for such a K that $\text{LCT}(K)$ is in PTIME and is $\text{L}_{\infty\omega}^\omega$ -axiomatizable, but not FO(LFP) -axiomatizable. \square

Exercise 8.4.16 ($\text{FO(PFP}_{\text{PTIME}}) \equiv \text{FO(LFP)} \Rightarrow \text{PSPACE} = \text{PTIME}$) Let $\tau_1 = \{E, U\}$ and $\tau_2 = \{E, U, V\}$, where U and V are unary and E is binary. Fix $k \geq 1$. Recall the definition of strings (= structures in STG) and

of labeled complete binary trees of the preceding exercise. For the present purpose we encode strings of length n in a suitable complete binary tree of depth n^k . The class TREE_k consists of τ_2 -structures $\mathcal{T} = (T, E^T, U^T, V^T)$ that satisfy

$$\begin{aligned} & (T, E^T, U^T) \text{ and } (T, E^T, V^T) \text{ are labeled complete binary trees;} \\ & \forall x \forall y ((Exy \wedge Vy) \rightarrow Vx) \\ & \forall x (Ux \rightarrow Vx) \end{aligned}$$

“The depth of \mathcal{T} is m^k , where m is the number of levels labeled by V ”.

TREE_k is axiomatizable in $\text{FO}(\text{LFP})$; in particular, the last condition can be expressed in $\text{FO}(\text{LFP})$ by developing the arithmetic of Example 8.1.11 for complete binary trees instead of orderings, the j -th level taking over the role of the j -th element of the ordering.

Consider the relation \asymp between strings and structures in TREE_k defined by using \approx of the preceding exercise as follows:

$$\mathcal{A} \asymp \mathcal{T} \quad \text{iff} \quad \mathcal{A} \approx (V^T, E^T \cap (V^T \times V^T), U^T).$$

Show:

- (a) If $K \subseteq \text{STG}$ is axiomatizable in $\text{FO}(\text{PFP})$, then for some k , $\{\mathcal{T} \in \text{TREE}_k \mid \mathcal{A} \asymp \mathcal{T} \text{ for some } \mathcal{A} \in K\}$ is axiomatizable in $\text{FO}(\text{PFP}_{\text{PTIME}})$.
- (b) If $H \subseteq \text{TREE}_k$ is axiomatizable in $\text{FO}(\text{LFP})$, then $\{\mathcal{A} \in \text{STG} \mid \mathcal{A} \asymp \mathcal{T} \text{ for some } \mathcal{T} \in H\}$ is axiomatizable in $\text{FO}(\text{LFP})$.

Let $K \subseteq \text{STG}$ be in PSPACE and therefore, axiomatizable in $\text{FO}(\text{PFP})$. Using (a), choose k such that $\{\mathcal{T} \in \text{TREE}_k \mid \mathcal{A} \asymp \mathcal{T} \text{ for some } \mathcal{A} \in K\}$ is axiomatizable in $\text{FO}(\text{PFP}_{\text{PTIME}})$ and hence, in $\text{FO}(\text{LFP})$, if we assume $\text{FO}(\text{PFP}_{\text{PTIME}}) \equiv \text{FO}(\text{LFP})$. Then, by (b), K is axiomatizable in $\text{FO}(\text{LFP})$ and therefore, K is in PTIME . Hence, $\text{PSPACE} = \text{PTIME}$ by the remark following 7.5.7. \square

8.4.2 Fixed-Point Logic with Counting

We have seen that $\text{FO}(\text{LFP})$ does not strongly capture PTIME and remarked that it is still open whether there exists a logic strongly capturing PTIME (cf. the remarks after 7.5.19). While pursuing this problem, various extensions of $\text{FO}(\text{LFP})$ have been studied as candidates for strongly capturing PTIME . As an example we present fixed-point logic with counting.

Recall that, when proving Theorem 8.4.14, we used the fact that, on ordered structures, $\text{FO}(\text{LFP})$ can describe the arithmetics needed to code polynomially many stages of a fixed-point operation in one relation. But arithmetics is not all orderings allows to do: Using orderings we can enumerate the elements of a structure and thus successively look at each single element. For instance, by enumeration we can calculate in $\text{FO}(\text{LFP})$ the sizes

of definable subsets. In fact, let $\varphi(x)$ be an FO(LFP)-formula. Then for any ordered structure \mathcal{A} (with $\mathcal{A} \models \forall x\varphi(x)$),

$$a \text{ is the } \|\varphi^{\mathcal{A}}(_)\| \text{-th element of } <^{\mathcal{A}} \text{ iff } \mathcal{A} \models [\text{LFP}_{xu, X\chi}] \max u[a],$$

where

$$\begin{aligned} \chi(x, u, X) := & (x = \min \wedge \neg\varphi(x) \wedge u = \min) \vee (x = \min \wedge \varphi(x) \wedge S \min u) \\ & \vee \exists x'u'(Xx'u' \wedge Sx'u' \wedge ((\neg\varphi(x) \wedge u = u') \vee (\varphi(x) \wedge Su'u))). \end{aligned}$$

Fixed-point logic with counting, FO(IFP, #), is an extension of fixed-point logic that contains arithmetic (in the size of the structure) and allows to calculate cardinalities. If we only consider ordered structures, every query definable in fixed-point logic with counting is in PTIME; so in this case the logic FO(IFP, #) coincides with least fixed-point logic and hence, captures PTIME. On the other hand, FO(IFP, #) does not strongly capture PTIME. In fact, we shall see that FO(IFP, #) is (equivalent to) a fragment of $C_{\infty\omega}^{\omega}$, infinitary logic with counting quantifiers and finitely many variables, and that there is a graph property in PTIME which is not even expressible in this logic (cf. section 12.2). Nevertheless, FO(IFP, #) has turned out to be quite an important logic, since it “captures PTIME on many concrete classes” (cf. Chapter 11). Since FO(IFP, #) contains arithmetics (in the size of the structure), those results for FO(LFP) and ordered structures where we used the ordering only for arithmetical operations or size calculations translate to fixed-point logic with counting and arbitrary finite structures (compare 8.4.19 below and 9.2.3 in the next chapter).

In *Fixed-Point Logic with Counting*, FO(IFP, #), any τ -structure \mathcal{A} is thought of to be accompanied by its *number part* consisting of a disjoint virtual universe $\{0, 1, \dots, \|\mathcal{A}\|\}$ together with the natural ordering, the successor relation, and constants for the first and the last element, respectively. We assume that $\tau \cap \{<, S, \min, \max\} = \emptyset$. \mathcal{A} is viewed as a two-sorted structure, the elements of \mathcal{A} itself being of the first sort, the *point sort*, those of the virtual universe being of the second sort, the *number sort*. On the syntactic side we have variables x, y, z, \dots for the first sort (*point variables*) and variables μ, ν, \dots for the second sort (*number variables*). The formulas of FO(IFP, #) are obtained from atomic formulas by closing under the connectives, first-order quantification of both sorts, and inflationary fixed-points of mixed sorted relations; moreover, for any formula φ and variable x there is a term $\#_x\varphi$ of the second sort denoting the cardinality of the set of elements x satisfying φ . More precisely, terms of both sorts and formulas are given simultaneously by the following clauses:

- All first-order τ -terms (with variables x, y, \dots) are terms of the first sort; all first-order terms of vocabulary $\{<, S, \min, \max\}$ (with variables μ, ν, \dots) are terms of the second sort.
- All atomic τ -formulas and all atomic $\{<, S, \min, \max\}$ -formulas (with terms of the corresponding sort) are formulas.

- If X is a second-order variable of mixed arity n_1, n_2 then $X\bar{t}\bar{\rho}$ is a formula, where \bar{t} and $\bar{\rho}$ are sequences of terms of the first and second sort of lengths n_1 and n_2 , respectively.
- If φ and ψ are formulas then so are $\neg\varphi$ and $(\varphi \vee \psi)$.
- If φ is a formula then $\exists x\varphi$ and $\exists\nu\varphi$ are formulas.
- If φ is a formula then $[\text{IFP}_{\bar{x}\bar{\mu}, X}\varphi]\bar{t}\bar{\rho}$ is a formula (the lengths of $\bar{x}, \bar{\mu}, \bar{t}, \bar{\rho}$ are according to the mixed arity of X).
- If φ is a formula then $\#_x\varphi$ is a term of the second sort.

The variable x does not occur free in $\#_x\varphi$. The semantics should be clear: If $\varphi = \varphi(x, \bar{y})$ and $\bar{a} \in \mathcal{A}$, the value $\#_x\varphi(x, \bar{a})^{\mathcal{A}}$ is given by the number $\|\{a \in A \mid \mathcal{A} \models \varphi[a, \bar{a}]\}\|$.

Examples 8.4.17 (a) Let $\tau = \{P\}$ with unary P . Then $\#_x Px = \#_x \neg Px$ is a sentence of $\text{FO}(\text{IFP}, \#)$ not equivalent to any $L_{\infty\omega}^\omega$ -sentence (cf. 3.3.25(b) and 2.3.12).

(b) The class EVEN of structures of even cardinality is axiomatizable in $\text{FO}(\text{IFP}, \#)$ by

$$\neg[\text{IFP}_{\mu, X} (\mu = \min \vee \exists\nu (X\nu \wedge \text{“}\mu = \nu + 2\text{”}))] \max.$$

(c) For equivalence relations E_1 and E_2 the following sentence of $\text{FO}(\text{IFP}, \#)$ expresses that, for all cardinalities l , E_1 and E_2 have the same number of equivalence classes of cardinality l :

$$\forall\mu(\#_x(\#_y E_1 xy = \mu) = \#_x(\#_y E_2 xy = \mu)).$$

(d) For graphs the following formula of $\text{FO}(\text{IFP}, \#)$ says that the distance between x and y is λ :

$$[\text{IFP}_{z\mu, U} (z = x \wedge \mu = 0) \vee (\neg\exists\nu U z\nu \wedge \exists z'\exists\mu'(U z'\mu' \wedge E z'z \wedge S\mu'\mu))] y\lambda.$$

(e) Note that the negation symbol is superfluous in $\text{FO}(\text{IFP}, \#)$. In fact, if x is a variable not occurring in φ then

$$\models_{\text{fin}} \neg\varphi \leftrightarrow \#_x\varphi = \min. \quad \square$$

It should be clear how $\text{FO}(\text{PFP}, \#)$, *Partial Fixed-Point Logic with Counting*, is defined.

Proposition 8.4.18 $\text{FO}(\text{PFP}, \#) \leq C_{\infty\omega}^\omega$.

Proof. Given a sentence φ of $\text{FO}(\text{PFP}, \#)$ we show that it is equivalent to a sentence of the form

$$\bigvee_{n \geq 1} (\exists^=^n x x = x \wedge \varphi_n)$$

where φ_n is a sentence of first-order logic with counting quantifiers that captures the meaning of φ in models of cardinality n . The translation from

φ to φ_n , though easy in principle, is involved if one wants to give a clean definition (the reader is encouraged to translate the formulas in the preceding examples to $C_{\infty\omega}^w$). We give a translation in two steps: First, we pass from φ to a sentence φ'_n eliminating $\#$ and PFP, but using atomic formulas and first-order quantifications of the second sort (they will be eliminated in the second step). Let φ'_n result from φ by replacing any subformula $\#_x\varphi = \mu$ by

$$\bigvee_{i \leq n} (\exists \bar{m}^i x \varphi \wedge \text{“}\mu \text{ is the } i\text{-th element of } \bar{m}^i \text{”})$$

and by treating PFP operations as in the proof of 8.4.2, taking into account that we can replace the infinite disjunctions used there by finite ones, since we only consider structures of cardinality n . In the second step we pass to an FO(C)-sentence φ_n . More generally, for every formula $\psi(\bar{x}, \mu_1, \dots, \mu_k)$, where atomic subformulas and quantifications of the second sort are still allowed, and for every $\bar{m} = m_1 \dots m_k$ with $m_1, \dots, m_k \leq n$ we define an FO(C)-formula $\psi_{\bar{m}}(\bar{x})$ such that, in structures of cardinality n , $\psi_{\bar{m}}(\bar{x})$ says the same as $\psi(\bar{x}, m_1, \dots, m_k)$. We then set $\varphi_n := (\varphi'_n)_{\bar{m}}$ with $\bar{m} = \emptyset$ (note that φ'_n is a sentence). The main steps in an inductive definition of $\psi_{\bar{m}}(\bar{x})$ are

- If $\psi(\bar{x}, \bar{\mu}) = \mu_i < \mu_j$ then $\psi_{\bar{m}}(\bar{x}) := \begin{cases} \text{T} & \text{if } m_i < m_j \\ \text{F} & \text{otherwise} \end{cases}$
- If $\psi(\bar{x}, \bar{\mu}) = \exists \mu \chi(\bar{x}, \bar{\mu}, \mu)$ then $\psi_{\bar{m}}(\bar{x}) := \bigvee_{m \leq n} \chi_{\bar{m}m}(\bar{x})$. □

Since in the proof of 8.4.14 we only used arithmetical properties of the ordering, in essentially the same way we get:

Proposition 8.4.19 $\text{FO}(\text{IFP}, \#) \equiv \text{FO}(\text{PFP}_{\text{PTIME}}, \#)$. □

Here, $\text{FO}(\text{PFP}_{\text{PTIME}}, \#)$ denotes the fragment of $\text{FO}(\text{PFP}, \#)$ consisting of those formulas for which each PFP operation closes in polynomially many steps.

Notes 8.4.20 Theorem 8.4.5 goes back to [12]. The main result 8.4.10 is due to Abiteboul and Vianu [4]. The proof presented here is based on [30]. Fixed-point logic with counting originates in ideas of Immerman [92]; in the form described here it has been shaped by Grädel and Otto [57, 124]. Otto has extended Theorem 8.4.10 to fixed-point logic with counting; cf. [125].

8.5 Fixed-Point Logics and Second-Order Logic

The results of Chapter 7 give us some information about the relationship between fixed-point logics and second-order logic SO: Let $K = \text{Mod}(\varphi)$ for some sentence φ of $\text{FO}(\text{LFP})$. Then, $K_{<} \in \text{PTIME} \subseteq \text{NPTIME}$ and hence, $K = \text{Mod}(\psi)$ for some Σ_1^1 -sentence ψ (by 7.5.14). Therefore, we have

$$(1) \quad \text{FO(LFP)} \leq \Sigma_1^1.$$

As the class EVEN of finite sets of even cardinality is not axiomatizable in FO(LFP) (see 8.4.4) but in Σ_1^1 (see the beginning of section 3.1), we have $\Sigma_1^1 \not\leq \text{FO(LFP)}$. On ordered structures the statement $\Sigma_1^1 \not\leq \text{FO(LFP)}$ is equivalent to $\text{NPTIME} \neq \text{PTIME}$ (cf. 7.5.7) and hence, is still open.

Since every class of ordered structures axiomatizable in SO is in PSPACE (cf. 7.4.3(b)), the Main Theorem (cf. 7.5.2) yields

$$(2) \quad \text{SO} \leq \text{FO(PFP)} \text{ on ordered structures.}$$

In this section we aim at purely modeltheoretic proofs of (1) and (2). The proof of (1) leads to an Ehrenfeucht-Fraïssé method for FO(LFP) that can also be used to get nonaxiomatizability results for FO(LFP). (However, in most cases it is easier to prove nonaxiomatizability even for $L_{\infty\omega}^w$, as we did in the last section for EVEN.)

Theorem 8.5.1 *On ordered structures, $\text{SO} \leq \text{FO(PFP)}$ and $\text{SO(PFP)} \equiv \text{FO(PFP)}$.*¹²

Proof. Let τ be a vocabulary with $\{<, S, \min, \max\} \subseteq \tau$. We restrict our attention to finite structures whose universe is $\{0, \dots, n\}$ for some n and whose $\{<, S, \min, \max\}$ -reduct is the natural ordering on $\{0, \dots, n\}$. For every formula φ of SO (or, of SO(PFP)) we inductively define an equivalent FO(PFP)-formula, the only nontrivial case being second-order quantification. So let $\varphi = \forall Z\psi(Z)$ and, for simplicity, assume that Z is unary. Below we show the existence of a first-order formula $\chi(x, y, X)$ with binary X such that in any ordered structure \mathcal{A} the fixed-point F_∞^χ exists and is nonempty, and the sequence of sets $(\{x \mid \exists y F_k^\chi xy\})_{k \geq 0}$ runs through all subsets of A . Then $\forall Z\psi(Z)$ is equivalent to

$$\exists x \exists y [\text{PFP}_{xy, X} (\chi(x, y, X) \wedge \psi(\exists y X _y))] xy.$$

To define χ , we code a sequence (a_0, \dots, a_n) of elements of A by the set $X := \{(a_i, i) \mid i \leq n\}$. It is easy to write down a formula $\chi(x, y, X)$ such that the stages $F_1^\chi, F_2^\chi, \dots, F_{(n+1)n+1}^\chi$ are codes of the sequences

$$(0, \dots, 0), (0, \dots, 0, 1), \dots, (n, \dots, n),$$

respectively, and $F_\infty^\chi = F_{(n+1)n+1}^\chi$. □

We now show that $\text{FO(LFP)} \leq \Sigma_1^1$. For the least fixed-point F_∞ of a monotone operation $F : \text{Pow}(A^k) \rightarrow \text{Pow}(A^k)$ we have (see part (b) of 8.1.2)

$$F_\infty = \bigcap \{X \subseteq A^k \mid F(X) \subseteq X\}.$$

¹² Cf. 7.5.19

Thus

$$\begin{aligned} \bar{a} \in F_\infty & \quad \text{iff} & \quad \text{for all } X \subseteq A^k, \text{ if } F(X) \subseteq X \text{ then } \bar{a} \in X \\ & \quad \text{iff} & \quad \text{for all } X \subseteq A^k, \text{ if } \bar{a} \notin X \text{ then there is some} \\ & & \quad \bar{b} \text{ with } \bar{b} \notin X \text{ and } \bar{b} \in F(X). \end{aligned}$$

Hence, for an arbitrary formula φ of FO(LFP), we have

$$(*) \quad \models_{\text{fin}} [\text{LFP}_{\bar{x}, X} \varphi] \bar{t} \leftrightarrow \forall X (\neg X \bar{t} \rightarrow \exists \bar{x} (\neg X \bar{x} \wedge \varphi)).$$

Since, by 8.2.9, every formula of FO(LFP) is equivalent to a formula of the form $\forall v [\text{LFP}_{\bar{x}, X} \varphi] \bar{v}$ with first-order φ , (*) shows that $\text{FO(LFP)} \leq \Pi_1^1$. Since FO(LFP) is closed under negation, we get

Theorem 8.5.2 $\text{FO(LFP)} \leq \Delta_1^1$. □

The equivalence (*) can be used to obtain a characterization in the Ehrenfeucht-Fraïssé style of classes axiomatizable in FO(LFP). Essentially, the corresponding Ehrenfeucht-Fraïssé game is obtained by taking into account relation moves that correspond to the right side of (*). Suppose that in a play the spoiler aims at showing

$$\mathcal{A} \models [\text{LFP}_{x, X} \varphi(x, X)] y [a] \quad \text{and} \quad \mathcal{B} \not\models [\text{LFP}_{x, X} \varphi(x, X)] y [b]$$

(a and b being the interpretations of y), that is, by (*),

$$(+)$$

$$\begin{array}{l} \mathcal{A} \models \forall X (\neg X y \rightarrow \exists x (\neg X x \wedge \varphi)) [a] \\ \text{and } \mathcal{B} \not\models \forall X (\neg X y \rightarrow \exists x (\neg X x \wedge \varphi)) [b]. \end{array}$$

For this purpose he starts by choosing a subset B_0 of B such that $b \notin B_0$ and $\mathcal{B} \not\models \exists x (\neg X x \wedge \varphi(x, X)) [B_0]$. The duplicator then selects a subset A_0 of A with $a \notin A_0$ claiming that it behaves like B_0 . However, by (+),

$$\mathcal{A} \models \exists x (\neg X x \wedge \varphi(x, X)) [A_0].$$

Thus the spoiler can choose an element $a' \notin A_0$ such that $\mathcal{A} \models \varphi[a', A_0]$, while for any $b' \notin B_0$ chosen by the duplicator we have $\mathcal{B} \not\models \varphi[b', B_0]$. Now the spoiler has a “simpler” task, namely to show that $\mathcal{A} \models \varphi[a', A_0]$ and $\mathcal{B} \not\models \varphi[b', B_0]$.

The order and type of choices of the spoiler and the duplicator in this example will be the defining clauses of a so-called LFP move in the game we are going to introduce. To simplify the notation and the presentation we consider only relational vocabularies and *monadic least fixed-point logic* FO(M-LFP), which is obtained by restricting FO(LFP) to formulas in which only unary relation variables (set variables) are allowed. By (*), we have $\text{FO(M-LFP)} \leq \text{MSO}$. Furthermore, we remark that FO(M-LFP) like monadic second-order logic MSO has turned out to be an interesting logic in the realm of infinite

structures (cf. [14, 41]).

We write $\mathcal{A} \equiv_m^M \mathcal{B}$ iff \mathcal{A} and \mathcal{B} satisfy the same sentences of FO(M-LFP) of quantifier rank $\leq m$, where the quantifier rank is defined inductively as for first-order logic, adding the clause

$$\text{qr}([\text{LFP}_{x,X} \varphi] y) := 1 + \text{qr}(\varphi).$$

In order to determine the winner of a play we need the notion of a generalized partial isomorphism. It takes into consideration that in $[\text{LFP}_{x,X} \varphi] y$ the formula φ is positive in X .

Definition 8.5.3 Let \mathcal{A} and \mathcal{B} be τ -structures (where τ is relational). Suppose that $\bar{a} = a_1 \dots a_k \in A$, $\bar{b} = b_1 \dots b_k \in B$, that $\bar{U} = U_1 \dots U_l$ and $\bar{V} = V_1 \dots V_n$ are sequences of subsets of A , and that $\bar{S} = S_1 \dots S_l$ and $\bar{T} = T_1 \dots T_n$ are sequences of subsets of B . Then $(\bar{a} \mapsto \bar{b}, \bar{U} \mapsto \bar{S}, \bar{V} \mapsto \bar{T})$ is a *generalized partial isomorphism* from \mathcal{A} to \mathcal{B} iff $\bar{a} \mapsto \bar{b}$ is a partial isomorphism from \mathcal{A} to \mathcal{B} and for $i = 1, \dots, k$ and $j = 1, \dots, l$ we have:

- $a_i \in U_j$ implies $b_i \in S_j$,
- $b_i \in T_j$ implies $a_i \in V_j$. □

As the Ehrenfeucht-Fraïssé game $G_m(\mathcal{A}, \mathcal{B})$ for first-order logic, the game MLFP- $G_m(\mathcal{A}, \mathcal{B})$ is played by two players, the spoiler and the duplicator, and consists of m moves. In the i -th move ($1 \leq i \leq m$) an element a_i of A and an element b_i of B are chosen, and in some moves subsets will be chosen, too. More precisely: There are two types of moves, *point moves* and *LFP moves*. In each move of a play the spoiler first decides what type of move he wants. Point moves are as in the first-order case. Now suppose that the spoiler decides that the i -th move should be an LFP move. Then he selects a $j < i$ and a structure, \mathcal{A} or \mathcal{B} . We speak of a *positive* LFP move if he selects \mathcal{A} , and of a *negative* LFP move if he selects \mathcal{B} .

In the positive case the spoiler chooses a subset S of B with $b_j \notin S$, and the duplicator answers by a subset U of A with $a_j \notin U$. Then the spoiler chooses an element $a_i \in A \setminus U$, and the move is finished by the selection of an element $b_i \in B \setminus S$ by the duplicator.

In the negative case the roles of A and B are interchanged, that is: The spoiler chooses a $V \subseteq A$ with $a_j \notin V$, and the duplicator answers by a $T \subseteq B$ with $b_j \notin T$; then the spoiler chooses an element $b_i \in B \setminus T$, and finally the duplicator selects an element $a_i \in A \setminus V$.

In the course of a play, besides the elements $a_1 \dots a_m$ and $b_1 \dots b_m$, subsets

$$U_1 \dots U_l \text{ of } A \quad \text{and} \quad S_1 \dots S_l \text{ of } B$$

have been chosen in positive LFP moves, and subsets

$$V_1 \dots V_n \text{ of } A \quad \text{and} \quad T_1 \dots T_n \text{ of } B$$

in negative ones. By definition, the duplicator wins the play if $(\bar{a} \mapsto \bar{b}, \bar{U} \mapsto \bar{S}, \bar{V} \mapsto \bar{T})$ is a generalized partial isomorphism.

The main theorem now reads as follows:

Theorem 8.5.4 *Let \mathcal{A} and \mathcal{B} be structures and $m \in \mathbb{N}$. Then the following are equivalent:*

- (i) $\mathcal{A} \equiv_m^M \mathcal{B}$, i.e., \mathcal{A} and \mathcal{B} satisfy the same sentences of FO(M-LFP) of quantifier rank $\leq m$.
- (ii) *The duplicator has a winning strategy in the game MLFP-G $_m(\mathcal{A}, \mathcal{B})$.*

Before presenting the proof we give an example.

Example 8.5.5 Let \mathcal{G}_l and $\mathcal{G}_l \dot{\cup} \mathcal{G}_l$ be graphs consisting of one and two cycles of length $l + 1$, respectively. To be definite, set

$$\mathcal{G}_l := \{0, \dots, l\}, \quad E^{\mathcal{G}_l} := \{(i, i + 1) \mid i < l\} \cup \{(i + 1, i) \mid i < l\} \cup \{(0, l), (l, 0)\}$$

and $\mathcal{G}_l \dot{\cup} \mathcal{G}_l := \{0, \dots, l\} \times \{0, 1\}$. – The spoiler has a winning strategy in the game MLFP-G $_4(\mathcal{G}_l, \mathcal{G}_l \dot{\cup} \mathcal{G}_l)$, which is exemplified by the following table where we have underlined his selections. He plays according to the fact that

$$\begin{aligned} \mathcal{G}_l &\models \exists x \forall y [\text{LFP}_{z,Z} z = x \vee \exists u (Zu \wedge Euz)] y \\ \mathcal{G}_l \dot{\cup} \mathcal{G}_l &\not\models \exists x \forall y [\text{LFP}_{z,Z} z = x \vee \exists u (Zu \wedge Euz)] y. \end{aligned}$$

move	\mathcal{G}_l	$\mathcal{G}_l \dot{\cup} \mathcal{G}_l$	$j < i$
1	<u>0</u>	(b, k)	
2	a	<u>$(b, 1 - k)$</u>	
3	U (with $a \notin U$) <u>$a' := \min\{u \mid u \notin U\}$</u>	<u>$\{(c, k) \mid c \leq l\}$</u> $(d, 1 - k)$	2
4	<u>$a' - 1$</u>	?	

No fourth move of the duplicator will lead to a generalized partial isomorphism, because $a' - 1$ is connected with a' by an edge and $a' - 1 \in U$, but there is no element in $\mathcal{G}_l \dot{\cup} \mathcal{G}_l$ related to $(d, 1 - k)$ and $\{(c, k) \mid c \leq l\}$ in this way. (If $a' = 0$ then already in the third move there is no answer for the duplicator leading to a generalized partial isomorphism.) \square

Proof (of 8.5.4). First suppose that the duplicator has a winning strategy for MLFP-G $_m(\mathcal{A}, \mathcal{B})$. We have to show that $(\mathcal{A} \models \varphi \text{ iff } \mathcal{B} \models \varphi)$ holds for all FO(M-LFP)-sentences of quantifier rank $\leq m$. Essentially, the proof runs in the standard way: As usual, point moves take care of first-order quantifiers and LFP moves of LFP operators as outlined before introducing the game. More precisely: Suppose that in a play, where the duplicator has followed his

winning strategy, $\bar{a} \mapsto \bar{b}$, $\bar{U} \mapsto \bar{S}$ and $\bar{V} \leftarrow \bar{T}$ have been the choices in the first i moves; then show that

$$\mathcal{A} \models \psi(\bar{x}, \bar{X}, \bar{Y}) [\bar{a}, \bar{U}, \bar{V}] \quad \text{implies} \quad \mathcal{B} \models \psi(\bar{x}, \bar{X}, \bar{Y}) [\bar{b}, \bar{S}, \bar{T}]$$

for any FO(M-LFP)-formula ψ of quantifier rank $\leq m-i$ which is positive in \bar{X} and negative in \bar{Y} (a corresponding inductive proof runs on $m-i$).

For the other direction we introduce an analogue of j -isomorphism types (cf. 2.2.5). Fix a structure \mathcal{A} . For $\bar{a} = a_1 \dots a_k \in A$, “positive” subsets $\bar{U} = U_1 \dots U_l$ of A and “negative” subsets $\bar{V} = V_1 \dots V_n$ of A let

$$\chi_{\bar{a}, \bar{U}, \bar{V}}^0 := \varphi_{\bar{a}}^0 \wedge \bigwedge_{1 \leq i \leq k} \left(\bigwedge_{1 \leq j \leq l; a_i \in U_j} X_j v_i \wedge \bigwedge_{1 \leq j \leq n; a_i \notin V_j} \neg Y_j v_i \right),$$

where $\varphi_{\bar{a}}^0$ is the (formula denoting the) 0-isomorphism type of \bar{a} in \mathcal{A} . Then, for arbitrary \mathcal{B} , \bar{b} , \bar{S} , and \bar{T} , we have

$$\mathcal{B} \models \chi_{\bar{a}, \bar{U}, \bar{V}}^0 [\bar{b}, \bar{S}, \bar{T}] \quad \text{iff} \quad (\bar{a} \mapsto \bar{b}, \bar{U} \mapsto \bar{S}, \bar{V} \leftarrow \bar{T}) \text{ is a generalized partial isomorphism.}$$

For $j > 0$ we define the formula $\chi_{\bar{a}, \bar{U}, \bar{V}}^j$ such that $\mathcal{B} \models \chi_{\bar{a}, \bar{U}, \bar{V}}^j [\bar{b}, \bar{S}, \bar{T}]$ whenever in a play of the FO(M-LFP) game, in which $\bar{a} \mapsto \bar{b}$, $\bar{U} \mapsto \bar{S}$, and $\bar{V} \leftarrow \bar{T}$ have been selected so far, there is a strategy for the duplicator which allows him to answer j further moves correctly. Thus we set (writing v , X , and Y for v_{k+1} , X_{l+1} , and Y_{n+1} , respectively):

$$\begin{aligned} \chi_{\bar{a}, \bar{U}, \bar{V}}^j := & \bigwedge_{a \in A} \exists v \chi_{\bar{a}a, \bar{U}, \bar{V}}^{j-1} \wedge \forall v \bigvee_{a \in A} \chi_{\bar{a}a, \bar{U}, \bar{V}}^{j-1} \wedge \\ & \bigwedge_{1 \leq i \leq k} (\forall X (\neg X v_i \rightarrow \bigvee_{U \subseteq A; a_i \notin U} \bigwedge_{a \notin U} \exists v (\neg X v \wedge \chi_{\bar{a}a, \bar{U}, \bar{V}}^{j-1}))) \\ & \wedge \bigwedge_{V \subseteq A; a_i \notin V} \exists Y (\neg Y v_i \wedge \forall v (\neg Y v \rightarrow \bigvee_{a \notin V} \chi_{\bar{a}a, \bar{U}, \bar{V}}^{j-1}))). \end{aligned}$$

If $\bar{a}, \bar{U}, \bar{V}$ are empty, we write $\chi_{\mathcal{A}}^j$ for the sentence $\chi_{\bar{a}, \bar{U}, \bar{V}}^j$. We show that, up to equivalence,

$$(1) \quad \chi_{\bar{a}, \bar{U}, \bar{V}}^j \text{ is an FO(M-LFP)-formula of quantifier rank } \leq j.$$

This proves the claim: In fact, since $\mathcal{A} \models \chi_{\mathcal{A}}^m$, we get $\mathcal{B} \models \chi_{\mathcal{A}}^m$ from $\mathcal{A} \equiv_m^M \mathcal{B}$ and hence, the duplicator has a winning strategy in $\text{MLFP-G}_m(\mathcal{A}, \mathcal{B})$.

To obtain (1) we first note that a simple induction on j shows that, given the lengths of the sequences which appear as subscripts, there are only finitely many formulas of the form $\chi_{\bar{a}, \bar{U}, \bar{V}}^j$. Thus the disjunctions and conjunctions in the definition of $\chi_{\bar{a}, \bar{U}, \bar{V}}^j$ are finite.

The proof of (1) also proceeds by induction on j . In $\chi_{\bar{a}, \bar{U}, \bar{V}}^j$ the conjunct

$$\exists Y (\neg Y v_i \wedge \forall v (\neg Y v \rightarrow \bigvee_{a \notin V} \chi_{\bar{a}a, \bar{U}, \bar{V}V}^{j-1}))$$

is equivalent to

$$\neg \forall Y (\neg Y v_i \rightarrow \exists v (\neg Y v \wedge \neg \bigvee_{a \notin V} \chi_{\bar{a}a, \bar{U}, \bar{V}V}^{j-1})),$$

which, by (*) before Theorem 8.5.2, is equivalent to

$$\neg [\text{LFP}_{v, Y} \neg \bigvee_{a \notin V} \chi_{\bar{a}a, \bar{U}, \bar{V}V}^{j-1}] v_i$$

and hence, by induction hypothesis, to an FO(M-LFP)-formula of quantifier rank $\leq j$. Finally, we show how to rewrite

$$\forall X (\neg X v_i \rightarrow \bigvee_{U \subseteq A; a_i \notin U} \bigwedge_{a \notin U} \exists v (\neg X v \wedge \chi_{\bar{a}a, \bar{U}U, \bar{V}}^{j-1}))$$

as an FO(M-LFP)-formula of quantifier rank $\leq j$. This formula has the form

$$\forall X (\neg X v_i \rightarrow \bigvee_{\gamma \in \Gamma} \bigwedge_{\delta \in \Delta_\gamma} \exists v (\neg X v \wedge \varphi_{\gamma\delta}))$$

for suitable finite sets Γ and Δ_γ . Let Π be the set of functions h with domain Γ such that $h(\gamma) \in \Delta_\gamma$ for $\gamma \in \Gamma$. Then, as usual, we can interchange the order of \forall and \bigwedge , thus obtaining the equivalent formula

$$\forall X (\neg X v_i \rightarrow \bigwedge_{h \in \Pi} \bigvee_{\gamma \in \Gamma} \exists v (\neg X v \wedge \varphi_{\gamma h(\gamma)}))$$

which in turn is equivalent to

$$\bigwedge_{h \in \Pi} \forall X (\neg X v_i \rightarrow \exists v (\neg X v \wedge \bigvee_{\gamma \in \Gamma} \varphi_{\gamma h(\gamma)})).$$

But, using (*) again, this formula can be rewritten as

$$\bigwedge_{h \in \Pi} [\text{LFP}_{v, X} \bigvee_{\gamma \in \Gamma} \varphi_{\gamma h(\gamma)}] v_i.$$

Note that the $\varphi_{\gamma h(\gamma)}$ have quantifier rank $\leq j - 1$. □

The following corollaries are now proven as in the case of first-order logic.

Corollary 8.5.6 *For every FO(M-LFP)-sentence φ of quantifier rank m ,*

$$\models_{\text{fin}} \varphi \leftrightarrow \bigvee \{ \chi_{\mathcal{A}}^m \mid \mathcal{A} \text{ is a finite model of } \varphi \}. \quad \square$$

Corollary 8.5.7 *For a class K of finite structures the following are equivalent:*

- (i) K is not axiomatizable in FO(M-LFP).
- (ii) For each $m \geq 1$ there are finite structures \mathcal{A} and \mathcal{B} such that

$$\mathcal{A} \in K, \mathcal{B} \notin K, \text{ and } \mathcal{A} \cong_m^M \mathcal{B}$$

(where $\mathcal{A} \cong_m^M \mathcal{B}$ means that the duplicator has a winning strategy in MLFP-G $_m(\mathcal{A}, \mathcal{B})$). \square

8.5.1 Digression: Implicit Definability

Let \mathcal{L} be a logic. An n -ary query \mathcal{Q} on the class of τ -structures is *implicitly \mathcal{L} -definable* if – with R a new n -ary relation symbol – there is an $\mathcal{L}[\tau \cup \{R\}]$ -sentence $\psi(R)$ such that

- (1) $\models_{\text{fin}} \exists^=1 X \psi(X)$;
- (2) for any τ -structure \mathcal{A} , $(\mathcal{A}, \mathcal{Q}(\mathcal{A})) \models \psi(R)$.

We then say that $\psi(R)$ (or, $\psi(X)$) *implicitly \mathcal{L} -defines \mathcal{Q}* . In section 3.5 we used the notion of implicit definability in a slightly different way; essentially, (1) was replaced by $\models_{\text{fin}} \exists^{\leq 1} X \psi(X)$. In order to distinguish both notions, one sometimes speaks of weak (in the present case) and strong implicit definability.

Clearly, any $\mathcal{L}[\tau \cup \{R\}]$ -sentence $\psi(R)$ such that

$$\models_{\text{fin}} \exists^=1 X \psi(X)$$

implicitly \mathcal{L} -defines a query. Note that for first-order $\psi(X)$ the formula $\exists^=1 X \psi(X)$ is a second-order formula.

We show that any formula of FO(LFP) is equivalent to a first-order formula that contains relations implicitly definable in FO:

Theorem 8.5.8 *Given an FO(LFP)-formula $\varphi(\bar{x})$, there are a first-order formula $\psi_0(X)$ and a first-order formula $\psi_1(\bar{x}, X)$ (where the free variables always are among the displayed ones) such that*

- (a) $\models_{\text{fin}} \exists^=1 X \psi_0(X)$
- (b) $\models_{\text{fin}} \forall X (\psi_0(X) \rightarrow (\varphi(\bar{x}) \leftrightarrow \psi_1(\bar{x}, X)))$.

Proof. For notational simplicity let φ be a sentence. By 8.2.9, we can assume that $\varphi = \exists u[\text{LFP}_{\bar{z}, Z} \chi(\bar{z}, Z)] \bar{u}$ with first-order χ . Let Z be k -ary. Recall the definition of the stage comparison relation \leq_X associated with χ (denoted by \leq_{F^X} in section 8.2):

$$\bar{y} \leq_X \bar{z} \quad \text{iff} \quad \text{for some } n \geq 1, \bar{y} \in F_n^X \text{ and } \bar{z} \in F_n^X \setminus F_{n-1}^X$$

(\bar{y} and \bar{z} range over k -tuples).

We show that \leq_χ is implicitly FO-definable by a formula $\psi_0(\leq)$. Then we are done, since $\bar{t} \in F_\infty^\chi$ is equivalent to $\bar{t} \leq_\chi \bar{t}$ and hence, we have

$$\models_{\text{fin}} \forall \leq (\psi_0(\leq) \rightarrow (\varphi \leftrightarrow \exists u \bar{u} \leq \bar{u})).$$

When defining $\psi_0(\leq)$ we use the abbreviations:

$$\begin{aligned} \bar{y} < \bar{z} & \text{ for } \bar{y} \leq \bar{z} \wedge \neg \bar{z} \leq \bar{y} \\ \text{pred}(\bar{z}) & \text{ for } \{\bar{y} \mid \bar{y} < \bar{z}\} \\ \text{field}(\leq) & \text{ for } \{\bar{z} \mid \exists \bar{y} (\bar{y} \leq \bar{z} \vee \bar{z} \leq \bar{y})\}. \end{aligned}$$

As $\psi_0(\leq)$ we take the conjunction of (1)-(4) where

- (1) “ \leq is a reflexive and transitive relation on $\text{field}(\leq)$ ” \wedge
 $\forall \bar{y} \in \text{field}(\leq) \forall \bar{z} \in \text{field}(\leq) (\bar{y} \leq \bar{z} \vee \bar{z} \leq \bar{y})$
- (2) $\forall \bar{z} (\bar{z} \leq \bar{z} \leftrightarrow \chi(\bar{z}, \text{field}(\leq)))$
- (3) $\forall \bar{z} ((\bar{z} \leq \bar{z} \wedge \neg \exists \bar{y} \bar{y} < \bar{z}) \leftrightarrow \chi(\bar{z}, \emptyset))$
- (4) $\forall \bar{y} \forall \bar{z} \in \text{field}(\leq) (\bar{y} \leq \bar{z} \leftrightarrow \chi(\bar{y}, \text{pred}(\bar{z})))$.

We prove that $\psi_0(\leq)$ implicitly defines \leq_χ : In a model of $\psi_0(\leq)$ let \preceq be the ordering induced by \leq on the classes of the equivalence relation \sim given by

$$\bar{y} \sim \bar{z} \quad \text{iff} \quad \bar{y} \leq \bar{z} \text{ and } \bar{z} \leq \bar{y}.$$

Then, an induction on $n \geq 1$ using (3) and (4), shows that for $\bar{y} \in \text{field}(\leq)$ the class of \bar{y} is the n -th element of \preceq iff $\bar{y} \in F_n \setminus F_{n-1}$.

Together with (2) this yields that $\psi_0(\leq)$ implicitly defines \leq_χ . \square

Given a logic \mathcal{L} , let $\text{IMP}(\mathcal{L})$ be the logic that allows to define exactly those queries that are expressible in \mathcal{L} using implicitly \mathcal{L} -definable queries. More precisely: An $\text{IMP}(\mathcal{L})[\tau]$ -formula $\varphi(\bar{x})$ with free variables among \bar{x} is a tuple

$$(\psi_1(R_1), \dots, \psi_m(R_m), \psi(\bar{x}, R_1, \dots, R_m))$$

where $m \geq 0$, each $\psi_i(R_i)$ is an $\mathcal{L}[\tau \dot{\cup} \{R_i\}]$ -sentence, and $\psi(\bar{x}, R_1, \dots, R_m) \in \mathcal{L}[\tau \dot{\cup} \{R_1, \dots, R_m\}]$ has free variables among \bar{x} and

$$\models_{\text{fin}} \exists^{=1} X_1 \psi_1(X_1) \wedge \dots \wedge \exists^{=1} X_m \psi_m(X_m).$$

The meaning of $\varphi(\bar{x})$ is that of

$$\forall X_1 \dots \forall X_m ((\psi_1(X_1) \wedge \dots \wedge \psi_m(X_m)) \rightarrow \forall \bar{x} (\varphi(\bar{x}) \leftrightarrow \psi(\bar{x}, X_1, \dots, X_m))),$$

i.e., for any τ -structure \mathcal{A} and $\bar{a} \in A$,

$$\mathcal{A} \models \varphi[\bar{a}] \quad \text{iff} \quad \mathcal{A} \models \psi[\bar{a}, R_1, \dots, R_m]$$

where R_i is uniquely determined by $\mathcal{A} \models \psi_i[R_i]$. The preceding theorem shows

Corollary 8.5.9 $\text{FO(LFP)} \leq \text{IMP(FO)}$. □

It has been shown (cf. [29]) that $\text{IMP(FO)} \not\leq \text{FO(LPF)}$, even $\text{IMP(FO)} \not\leq L_{\infty\omega}^\omega$.

Since the formula

$$(\psi_1(R_1), \dots, \psi_m(R_m), \psi(\bar{x}, R_1, \dots, R_m))$$

in the definition of $\text{IMP}(\mathcal{L})$ is equivalent to

$$\forall X_1 \dots \forall X_m ((\psi_1(X_1) \wedge \dots \wedge \psi_m(X_m)) \rightarrow \psi(\bar{x}, X_1, \dots, X_m))$$

and to

$$\exists X_1 \dots \exists X_m (\psi_1(X_1) \wedge \dots \wedge \psi_m(X_m) \wedge \psi(\bar{x}, X_1, \dots, X_m))$$

we get

Proposition 8.5.10 $\text{IMP(FO)} \leq \Delta_1^1$. □

Exercise 8.5.11 Show: $\text{IMP(IMP(FO))} \equiv \text{IMP(FO)}$. □

Exercise 8.5.12 Show that $\text{IMP(FO(PFP))} \equiv \text{FO(PFP)}$ on ordered structures (use Theorem 8.5.1). □

Exercise 8.5.13 Let \mathcal{A} be a τ -structure. For an ordered representation $(\mathcal{A}, <)$ let $w_{(\mathcal{A}, <)}$ be the $\{0, 1\}$ -word obtained by concatenating the inscriptions of all input tapes of a Turing machine started with $(\mathcal{A}, <)$. For fixed \mathcal{A} , all orderings $<$ lead to words of the same lengths, say l .

(a) Show: $w_{(\mathcal{A}, <_1)} = w_{(\mathcal{A}, <_2)}$ iff $(\mathcal{A}, <_1) \cong (\mathcal{A}, <_2)$.

An ordering $<$ on \mathcal{A} is *minimal*, if for all orderings $<'$ on \mathcal{A} the word $w_{(\mathcal{A}, <')}$ is not smaller than $w_{(\mathcal{A}, <)}$ in the lexicographic ordering of $\{0, 1\}^l$. Hence, by (a), there is a unique minimal ordering iff \mathcal{A} is rigid. Let \mathcal{Q} be the query that assigns the minimal ordering to a rigid structure and the empty binary relation to a nonrigid one.

(b) Show that \mathcal{Q} is expressible in IMP(FO(PFP)) (and thus, $\text{IMP(FO(PFP))} \not\leq L_{\infty\omega}^\omega$ by 8.4.12). Hint: A formula $\psi(X)$ implicitly defines \mathcal{Q} if it expresses in any \mathcal{A} that “ X is an ordering, and for all orderings Y different from X we have: If π is the map with $\pi : (A, X) \cong (A, Y)$, then not $\pi : (\mathcal{A}, X) \cong (\mathcal{A}, Y)$, and if R is the first relation and \bar{a} the first X -tuple such that

$$R^{\mathcal{A}}\bar{a} \leftrightarrow R^{\mathcal{A}}\pi(\bar{a}) \text{ does not hold,}$$

then

$$\text{not } R^{\mathcal{A}}\bar{a} \text{ and } R^{\mathcal{A}}\pi(\bar{a}),$$

or there is no such ordering and $X = \emptyset$ ”. To get $\psi(X)$ as an FO(PFP) -formula note that in the presence of X and Y , the map π is definable in FO(LFP) and that, in the presence of the ordering X , the second-order quantification “for all Y ” can be expressed in FO(PFP) (cf. the proof of 8.5.1). □

Notes 8.5.14 The game-theoretic characterization of FO(LFP)-equivalence is due to Bosse [14]. Theorem 8.5.8 is due to Kolaitis [101], Exercise 8.5.13 is taken from [81].

8.6 Transitive Closure Logic

One of the most prominent examples of a global relation expressible in least fixed-point logic but not in first-order logic is the transitive closure of a given relation. *Transitive Closure Logic* FO(TC) is obtained by adding to first-order logic the operation of taking the transitive closure of definable relations. We have introduced FO(TC) already in Chapter 7, where we showed that FO(TC) captures NLOGSPACE. This section contains a more detailed analysis of the expressive power of FO(TC); it can be read independently of Chapter 7.

First we recall some definitions. Let $k \geq 1$ and X be a $2k$ -ary relation on a set A . The *transitive closure* $\text{TC}(X)$ of X is defined by

$$\text{TC}(X) := \{(\bar{a}, \bar{b}) \in A^{2k} \mid \text{there exist } n > 0 \text{ and } \bar{e}_0, \dots, \bar{e}_n \in A^k \text{ such that } \bar{a} = \bar{e}_0, \bar{b} = \bar{e}_n, \text{ and for all } i < n, (\bar{e}_i, \bar{e}_{i+1}) \in X\},$$

and the *deterministic transitive closure* $\text{DTC}(X)$ by

$$\text{DTC}(X) := \{(\bar{a}, \bar{b}) \in A^{2k} \mid \text{there exist } n > 0 \text{ and } \bar{e}_0, \dots, \bar{e}_n \in A^k \text{ such that } \bar{a} = \bar{e}_0, \bar{b} = \bar{e}_n, \text{ and for all } i < n, \bar{e}_{i+1} \text{ is the unique } \bar{e} \text{ for which } (\bar{e}_i, \bar{e}) \in X\}.$$

For a vocabulary τ the class of formulas of FO(TC) of vocabulary τ is given by the calculus with the following rules:

- (i) $\frac{}{\varphi}$ where φ is an atomic first-order formula over τ
- (ii) $\frac{\varphi}{\neg\varphi}, \frac{\varphi, \psi}{(\varphi \vee \psi)}, \frac{\varphi}{\exists x\varphi}$
- (iii) $\frac{\varphi}{[\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}}$ where the variables in \bar{x}, \bar{y} are pairwise distinct and where the tuples $\bar{x}, \bar{y}, \bar{s}$, and \bar{t} are all of the same length, \bar{s} and \bar{t} being tuples of terms.

For FO(DTC), rule (iii) is replaced by $\frac{\varphi}{[\text{DTC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}}$ with the same side conditions.

Sentences are formulas without free variables, where the free occurrence of a variable is defined in the standard way with the additional clause

$$\text{free}([\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}) := \text{free}(\bar{s}) \cup \text{free}(\bar{t}) \cup (\text{free}(\varphi) \setminus \{\bar{x}, \bar{y}\}),$$

and similarly for DTC.

The semantics is defined inductively w.r.t. the calculus above, the meaning of $[\text{TC}_{\bar{x},\bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s}\bar{t}$ being $(\bar{s}, \bar{t}) \in \text{TC}(\{(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{y})\})$; and $[\text{DTC}_{\bar{x},\bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s}\bar{t}$ expresses that $(\bar{s}, \bar{t}) \in \text{DTC}(\{(\bar{x}, \bar{y}) \mid \varphi(\bar{x}, \bar{y})\})$.¹³

Examples and Remarks 8.6.1 (a) $\text{FO}(\text{DTC}) \leq \text{FO}(\text{TC})$ since

$$\models_{\text{fin}} [\text{DTC}_{\bar{x},\bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s}\bar{t} \leftrightarrow [\text{TC}_{\bar{x},\bar{y}}(\varphi(\bar{x}, \bar{y}) \wedge \forall \bar{z}(\varphi(\bar{x}, \bar{z}) \rightarrow \bar{z} = \bar{y}))] \bar{s}\bar{t}.$$

(b) $\text{FO} \not\leq \text{FO}(\text{TC})$ since

$$\forall x \forall y (\neg x = y \rightarrow [\text{TC}_{x,y} Exy] xy)$$

expresses “connectivity” in the class of graphs (cf. 2.3.8).

(c) $\text{FO} \not\leq \text{FO}(\text{DTC})$ since

$$\forall x \forall y [\text{DTC}_{x,y} Exy] xy$$

is an $\text{FO}(\text{DTC})$ -sentence which, for every $l \geq 1$, holds in the digraph \mathcal{D}_l consisting of a cycle of length $l + 1$, but not in the union $\mathcal{D}_l \dot{\cup} \mathcal{D}_l$. Moreover,

$$\forall u \forall v \exists w \exists z [\text{DTC}_{xy, x'y'} (Exy \wedge Ex'y' \wedge x' = y \wedge \neg y' = x)] uwvz$$

separates the “undirected” cycles \mathcal{G}_l from $\mathcal{G}_l \dot{\cup} \mathcal{G}_l$ for $l \geq 2$.

(d) $\text{FO}(\text{TC}) \leq \text{FO}(\text{LFP})$ since

$$\models_{\text{fin}} [\text{TC}_{\bar{x},\bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s}\bar{t} \leftrightarrow [\text{LFP}_{\bar{x}\bar{y}, X} (\varphi(\bar{x}, \bar{y}) \vee \exists \bar{u} (X\bar{x}\bar{u} \wedge \varphi(\bar{u}, \bar{y})))] \bar{s}\bar{t}. \quad \square$$

Exercise 8.6.2 Show that the class of bipartite graphs is axiomatizable in $\text{FO}(\text{TC})$. Hint: Use 6.3.6. □

Exercise 8.6.3 For $r \geq 1$ let $\text{FO}(\text{TC}^r)$ be the fragment of $\text{FO}(\text{TC})$ consisting of those formulas, where for each subformula of the form $[\text{TC}_{\bar{x},\bar{y}} \varphi] \bar{s}\bar{t}$ we have $\text{length}(\bar{x}) = r$. Clearly, $\text{FO}(\text{TC}^1) \leq \text{FO}(\text{TC}^2) \leq \dots$. Show that $\text{FO}(\text{TC}^1) \equiv \text{MSO}$ on the class of word models (cf. section 6.2). (Hint: $\text{FO}(\text{TC}^1) \leq \text{MSO}$ is true in general; for $\text{MSO} \leq \text{FO}(\text{TC}^1)$ give a proof by induction on regular expressions (cf. section 6.1).)

Define $\text{FO}(\text{DTC}^r)$ similarly and show that $\text{FO}(\text{DTC}^2) \not\leq \text{MSO}$ on the class of word models. (Hint: Use the fact that the set $\{a^n b^n \mid n \geq 1\}$ is not regular.) □

8.6.1 $\text{FO}(\text{DTC}) < \text{FO}(\text{TC})$

By the results of Chapter 7 we know that

¹³ Recall that the notation $\varphi(\bar{x}, \bar{y})$ exhibits only the variables relevant in this situation. So φ may contain free variables that are not among \bar{x}, \bar{y} .

- $\text{FO}(\text{DTC}) \not\equiv \text{FO}(\text{TC})$ on ordered structures iff $\text{LOGSPACE} \neq \text{NLOGSPACE}$
- $\text{FO}(\text{TC}) \not\equiv \text{FO}(\text{LFP})$ on ordered structures iff $\text{NLOGSPACE} \neq \text{PTIME}$.

The inequalities on the right hand side are prominent open problems of complexity theory. The following remarks will lead to a proof of $\text{FO}(\text{DTC}) \not\equiv \text{FO}(\text{TC})$ for *arbitrary* finite structures; the corresponding inequality $\text{FO}(\text{TC}) \not\equiv \text{FO}(\text{LFP})$ will be shown in subsection 8.6.3.

Definition 8.6.4 A subset C of a structure \mathcal{A} is said to be *closed* if for all $a \in A \setminus C$ there is an automorphism π of \mathcal{A} pointwise fixing C with $\pi(a) \neq a$. □

Every DTC-path starting in a closed set must stay in it:

Lemma 8.6.5 *Let C be a closed subset of a structure \mathcal{A} . For every formula $\varphi(\bar{x}, \bar{y}, \bar{z})$ with free variables among $\bar{x}, \bar{y}, \bar{z}$ and for every $\bar{c} \in C$,*

$$\mathcal{A} \models [\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{c})] \bar{x} \bar{y} [\bar{a} \bar{b}] \text{ and } \bar{a} \in C \text{ imply } \bar{b} \in C.$$

Proof. Since $(\bar{a}, \bar{b}) \in \text{DTC}(\varphi^{\mathcal{A}}(-, -, \bar{c}))$, there is a sequence $\bar{e}_0, \dots, \bar{e}_m$ such that

$$(*) \quad \bar{a} = \bar{e}_0, \bar{b} = \bar{e}_m, \text{ and for } i < m, \mathcal{A} \models \forall \bar{y} (\varphi(\bar{e}_i, \bar{y}, \bar{c}) \leftrightarrow \bar{y} = \bar{e}_{i+1}).$$

If $\bar{b} \notin C$ then there is a $j < m$ such that $\bar{e}_j \in C$ and $\bar{e}_{j+1} \notin C$, say, $\bar{e}_{j+1} = e_1 \dots e_k$ and $e_i \notin C$. Choose an automorphism π of \mathcal{A} pointwise fixing C with $\pi(e_i) \neq e_i$. Then

$$\mathcal{A} \models \varphi[\bar{e}_j, \bar{e}_{j+1}, \bar{c}], \quad \mathcal{A} \models \varphi[\bar{e}_j, \pi(\bar{e}_{j+1}), \bar{c}], \quad \text{and } \bar{e}_{j+1} \neq \pi(\bar{e}_{j+1}),$$

contradicting (*). □

By the lemma there is a bound for the length of DTC-paths without repetitions that start in a closed set of a given cardinality. This is the reason why in classes of structures with many closed sets of controllable cardinality, so-called indiscernible classes, every DTC operation is first-order definable. We thus are led to a proof of $\text{FO}(\text{DTC}) \not\equiv \text{FO}(\text{TC})$ by showing

- $\text{FO}(\text{DTC}) \equiv \text{FO}$ on every indiscernible class.
- $\text{FO}(\text{TC}) \not\equiv \text{FO}$ on some indiscernible class.

Definition 8.6.6 Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be monotone (that is, $m \leq n$ implies $f(m) \leq f(n)$). A structure \mathcal{A} is said to be *f-indiscernible*, if for every subset X of A there is a closed set Y such that $X \subseteq Y \subseteq A$ and $\|Y\| \leq f(\|X\|)$. A class K of structures is *indiscernible* if, for some monotone function f , every structure in K is *f-indiscernible*. □

Proposition 8.6.7 *Let K be an indiscernible class. Then, on K , every formula of $\text{FO}(\text{DTC})$ is equivalent to a first-order formula.*

Proof. Let K be indiscernible via f and $\varphi(\bar{x}, \bar{y}, \bar{z})$ with $\bar{x} = x_1 \dots x_k$, $\bar{y} = y_1 \dots y_k$, and $\bar{z} = z_1 \dots z_l$ be an FO-formula. Suppose that $\mathcal{A} \in K$, $\bar{c} \in A$, and $(\bar{a}, \bar{b}) \in \text{DTC}(\varphi^{\mathcal{A}}(-, -, \bar{c}))$. Choose a closed set $C \subseteq A$ such that $\{\bar{a}, \bar{c}\} \subseteq C$ and $\|C\| \leq m := f(k+l)$. By the preceding lemma, every deterministic path connecting \bar{a} and \bar{b} stays within C ; hence, the shortest such path has length $\leq m^k$. Therefore, $[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}, \bar{z})] \bar{s} \bar{t}$ is equivalent to a first-order formula expressing “there exists a deterministic φ -path of length $\leq m^k$ from \bar{s} to \bar{t} ”. \square

For any structure \mathcal{A} in a relational vocabulary τ we denote, as in section 5.2, by $\mathcal{A} \times 2$ the structure we get by duplicating each element. To be precise, let $\mathcal{A} \times 2$ be the τ -structure with universe $A \times \{0, 1\}$ such that for any n -ary $R \in \tau$,

$$R^{A \times 2} := \{((a_1, i_1), \dots, (a_n, i_n)) \mid R^A a_1 \dots a_n \text{ and } i_1, \dots, i_n \in \{0, 1\}\}.$$

Lemma 8.6.8 $\mathcal{A} \times 2$ is f_0 -indiscernible, where $f_0(m) = 2 \cdot m$.

Proof. Clearly, for $X \subseteq A \times 2$ the set

$$C(X) := X \cup \{(a, i) \mid (a, 1-i) \in X\}$$

is closed: for any $(b, i) \in (A \times 2) \setminus C(X)$, the function $\pi_b : A \times 2 \rightarrow A \times 2$ given by

$$\pi_b((a, j)) := \begin{cases} (a, j) & \text{if } a \neq b \\ (b, 1-j) & \text{if } a = b \end{cases}$$

is the desired automorphism. \square

Theorem 8.6.9 $\text{FO}(\text{DTC}) \not\equiv \text{FO}(\text{TC})$.

Proof. First of all note that

$$(*) \quad \mathcal{A} \cong_m \mathcal{B} \text{ implies } \mathcal{A} \times 2 \cong_m \mathcal{B} \times 2.$$

Let $K := \{\mathcal{G}_l \mid l \geq 2\}$ be the set of cycles and let $H := \{\mathcal{G}_l \dot{\cup} \mathcal{G}_l \mid l \geq 2\}$. By the preceding lemma the sets $K \times 2$ ($:= \{\mathcal{A} \times 2 \mid \mathcal{A} \in K\}$) and $H \times 2$ are indiscernible. Since $\mathcal{G}_l \cong_m \mathcal{G}_l \dot{\cup} \mathcal{G}_l$ for $l \geq 2^m$ (cf. 2.3.8), $(*)$ shows that there is no first-order sentence separating $K \times 2$ and $H \times 2$, and hence by 8.6.7 and 8.6.8, no $\text{FO}(\text{DTC})$ -sentence. But the structures in $K \times 2$ are connected graphs and those in $H \times 2$ are not; thus there is a separating $\text{FO}(\text{TC})$ -sentence (cf. 8.6.1(b)). \square

Exercise 8.6.10 Show that the class of bipartite graphs is not axiomatizable in $\text{FO}(\text{DTC})$. \square

8.6.2 FO(posTC) and Normal Forms

Let FO(posDTC) and FO(posTC) be the logics consisting of the formulas of FO(DTC) and FO(TC) in which DTC and TC, respectively, only occur positively, i.e., within the scope of an even number of negation signs (we assume that the formulas of FO(DTC) are built up from atomic formulas using $\neg, \wedge, \vee, \forall, \exists$, and DTC; similarly for FO(TC)).

Proposition 8.6.11 $\text{FO}(\text{posDTC}) \equiv \text{FO}(\text{DTC})$.

Proof. The proof of the nontrivial part proceeds by induction on the number of occurrences of logical operations ($\neg, \wedge, \vee, \forall, \exists, \text{DTC}$) in a formula of FO(DTC). In the main step of the induction we have to show that an FO(DTC)-formula of the form

$$(*) \quad \neg[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}$$

is equivalent to an FO(posDTC)-formula where, by induction hypothesis, we may assume that φ and $\neg\varphi$ are equivalent to FO(posDTC)-formulas. Obviously, (*) just says that the deterministic φ -path starting at \bar{s} reaches, without passing through \bar{t} , a point \bar{z} , where it ends or where a φ -cycle starts that does not contain \bar{t} . Noting that the deterministic φ -path ends at \bar{z} if \bar{z} has no or more than one φ -neighbour, we thus can express $\neg[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}$ by

$$\begin{aligned} & \exists \bar{z} ((\bar{s} = \bar{z} \vee [\text{DTC}_{\bar{x}, \bar{y}} (\varphi(\bar{x}, \bar{y}) \wedge \neg \bar{y} = \bar{t})] \bar{s} \bar{z}) \wedge (\forall \bar{x} \neg \varphi(\bar{z}, \bar{x}) \\ & \vee \exists \bar{x} \exists \bar{y} (\varphi(\bar{z}, \bar{x}) \wedge \varphi(\bar{z}, \bar{y}) \wedge \neg \bar{x} = \bar{y}) \vee [\text{DTC}_{\bar{x}, \bar{y}} (\varphi(\bar{x}, \bar{y}) \wedge \neg \bar{y} = \bar{t})] \bar{z} \bar{z})). \end{aligned}$$

Applying the induction hypothesis to φ and $\neg\varphi$ we see that this formula is equivalent to an FO(posDTC)-formula. \square

Note that the maximal number of *nested* occurrences of the DTC operator in an FO(DTC)-formula and in the FO(posDTC)-formula assigned to it in the preceding proof are the same. We use this fact when showing

Proposition 8.6.12 $\text{FO}(\text{DTC}) \leq \text{FO}(\text{posTC})$.

Proof. By the preceding proposition it suffices to show that $\text{FO}(\text{posDTC}) \leq \text{FO}(\text{posTC})$. The proof is by induction on the number of nested DTC operators, the only nontrivial step being the case of an FO(DTC)-formula of the form

$$[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}.$$

By the remark preceding the proposition, $\neg\varphi$ is equivalent to an FO(posDTC)-formula with the same maximal number of nested occurrences of the DTC operator as φ . Thus, by induction hypothesis, both φ and $\neg\varphi$ are equivalent to formulas of FO(posTC); therefore the same holds for $[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}$, since this formula is equivalent to $[\text{TC}_{\bar{x}, \bar{y}} (\varphi(\bar{x}, \bar{y}) \wedge \forall \bar{z} (\neg \varphi(\bar{x}, \bar{z}) \vee \bar{z} = \bar{y}))] \bar{s} \bar{t}$. \square

We have already shown in Chapter 7 that $\text{FO}(\text{posTC}) \equiv \text{FO}(\text{TC})$ on ordered structures. Below, we give a purely modeltheoretic proof of this fact. On arbitrary structures, $\text{FO}(\text{posTC})$ has less expressive power than $\text{FO}(\text{TC})$. The next proposition contains a normal form for a fragment of $\text{FO}(\text{posTC})$ that will be useful for our purposes.

Proposition 8.6.13 *Suppose that τ contains two constants c and d . Let φ be an existential $\text{FO}(\text{posTC})$ -formula, that is a formula built up from atomic and negated atomic formulas with the help of \wedge , \vee , TC , and \exists . Then, in models of $\neg c = d$, φ is equivalent to a formula of the form*

$$[\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{c} \tilde{d}$$

where ψ is quantifier-free.

Proof. The proof proceeds by induction on φ . Of course, any first-order formula φ is equivalent to $[\text{TC}_{\bar{x}, \bar{x}'} \varphi] \tilde{c} \tilde{d}$, where \bar{x}, \bar{x}' are variables not occurring in φ . Let $\varphi = (\varphi_1 \vee \varphi_2)$ and suppose that φ_i is equivalent to $[\text{TC}_{\bar{x}, \bar{x}'} \psi_i(\bar{x}, \bar{x}')] \tilde{c} \tilde{d}$ with quantifier-free ψ_i (add dummy variables to obtain sequences \bar{x}, \bar{x}' of the same length for φ_1 and φ_2). Then φ will be equivalent to a formula $[\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{c} \tilde{d}$, where the additional arguments x and x' in ψ are used as a switch as indicated by the figure

$$\begin{array}{ccc} \tilde{c}c & \dots \xrightarrow{\psi_1} & \dots \tilde{d}c \\ \downarrow & & \downarrow \\ \tilde{c}d & \dots \xrightarrow{\psi_2} & \dots \tilde{d}d \end{array}$$

In the formula ψ below, the first line corresponds to the upper path in the figure leading from $\tilde{c}c$ to $\tilde{d}d$, and the second line to the lower path:

$$\begin{aligned} \psi(\bar{x}x, \bar{x}'x') & := (x = x' = c \wedge \psi_1) \vee (\bar{x} = \bar{x}' = \tilde{d} \wedge x = c \wedge x' = d) \\ & \vee (\bar{x} = \bar{x}' = \tilde{c} \wedge x = c \wedge x' = d) \vee (x = x' = d \wedge \psi_2). \end{aligned}$$

Similarly, for the conjunction $(\varphi_1 \wedge \varphi_2)$ we take as $\psi(\bar{x}x, \bar{x}'x')$ a formula expressing the existence of a path as indicated by

$$\tilde{c}c \dots \xrightarrow{\psi_1} \dots \tilde{d}c \longrightarrow \tilde{c}d \dots \xrightarrow{\psi_2} \dots \tilde{d}d.$$

Now suppose that $\varphi = [\text{TC}_{\bar{u}, \bar{v}} \varphi'] \bar{s} \bar{t}$ where, by induction hypothesis, we can assume that φ' has the form $[\text{TC}_{\bar{x}, \bar{x}'} \psi'(\bar{x}, \bar{x}', \bar{u}, \bar{v})] \tilde{c} \tilde{d}$ with quantifier-free ψ' . Let $\bar{s} = \bar{e}_0, \dots, \bar{e}_k = \bar{t}$ be a path witnessing that $(\bar{s}, \bar{t}) \in \text{TC}(\varphi'(-, -))$. Then, for $i = 0, \dots, k-1$ there is a ψ' -path witnessing that $(\tilde{c}, \tilde{d}) \in \text{TC}(\psi'(-, -, \bar{e}_i, \bar{e}_{i+1}))$. Therefore, φ is equivalent to

$$[\text{TC}_{\bar{u}, \bar{v}, \bar{x}, \bar{x}'} \psi(\bar{u}, \bar{v}, \bar{x}, \bar{x}')] \tilde{c} \tilde{c} \tilde{d} \tilde{d} \tilde{d}$$

where in the formula ψ given below the first line sets the corresponding starting values, the second line takes care of witnessing that $(\tilde{c}, \tilde{d}) \in$

$\text{TC}(\psi'(-, -, \bar{e}_i, \bar{e}_{i+1}))$, the next line allows to pass from the tuple \bar{e}_i, \bar{e}_{i+1} to $\bar{e}_{i+1}, \bar{e}_{i+2}$, and the last line realizes that $\bar{e}_k = \bar{t}$:

$$\begin{aligned} \psi(\bar{u}, \bar{v}, \bar{x}, \bar{u}', \bar{v}', \bar{x}') &:= (\bar{u} = \bar{c} \wedge \bar{v} = \bar{c} \wedge \bar{x} = \bar{c} \wedge \bar{u}' = \bar{s} \wedge \bar{x}' = \bar{c}) \\ &\vee (\neg \bar{x} = \bar{d} \wedge \bar{u}' = \bar{u} \wedge \bar{v}' = \bar{v} \wedge \psi'(\bar{x}, \bar{x}', \bar{u}, \bar{v})) \\ &\vee (\bar{x} = \bar{d} \wedge \neg \bar{v} = \bar{t} \wedge \bar{u}' = \bar{v} \wedge \bar{x}' = \bar{c}) \\ &\vee (\bar{x} = \bar{d} \wedge \bar{v} = \bar{t} \wedge \bar{u}' = \bar{d} \wedge \bar{v}' = \bar{d} \wedge \bar{x}' = \bar{d}). \end{aligned}$$

For the existential quantifier note that

$$\models_{\text{fin}} \exists z \varphi(z) \leftrightarrow [\text{TC}_{x,y}((x = c \wedge \varphi(y)) \vee (\varphi(x) \wedge y = d))] \text{cd}.$$

So we can apply what we already know to obtain the desired normal form also for this case. \square

Corollary 8.6.14 *Let $\tau \supseteq \{<, S, \min, \max\}$. On ordered τ -structures with at least two elements, every $\text{FO}(\text{posTC})$ -formula is equivalent to a formula of the form*

$$[\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \widetilde{\min \max}$$

with quantifier-free ψ .

Proof. In ordered models, $\forall x \varphi(x)$ is equivalent to

$$[\text{TC}_{x,y}(\varphi(x) \wedge Sxy \wedge \varphi(y))] \min \max.$$

Hence, the universal quantifier can be eliminated in $\text{FO}(\text{posTC})$ -formulas. Thus the claim follows from the preceding proposition. \square

We use this normal form for $\text{FO}(\text{posTC})$ -formulas when proving now that $\text{FO}(\text{posTC}) \equiv \text{FO}(\text{TC})$ on ordered structures.

Fix a vocabulary τ for ordered structures, for simplicity, $\tau \supseteq \{<, S, \min, \max\}$. We saw in 8.1.11 that for $k \geq 1$ the successor relation S^k of the lexicographic ordering of k -tuples is first-order definable. For any ordered τ -structure \mathcal{A} with $A = \{0, \dots, n-1\}$ and for any $m_1, \dots, m_k \in A$ we set

$$[m_1, \dots, m_k] := m_1 \cdot n^{k-1} + \dots + m_{k-1} \cdot n + m_k.$$

In the following we use self-explanatory notations such as $\bar{l} = \overline{m} + 1$ for $S^k \overline{m}$ or, equivalently, for $[\bar{l}] = [\overline{m}] + 1$. Often we write 0 for min.

Theorem 8.6.15 *On ordered structures, $\text{FO}(\text{posTC}) \equiv \text{FO}(\text{TC})$.*

Proof. The proof – a purely modeltheoretic version of the proof given in 7.5.20 – proceeds by induction on $\text{FO}(\text{TC})$ -formulas. The only nontrivial case is the negation step. By induction hypothesis and the corollary above it suffices to show for first-order ψ that

$$\neg[\text{TC}_{\bar{x}, \bar{y}} \psi(\bar{x}, \bar{y})] \bar{s} \bar{t}$$

is equivalent to an FO(posTC)-formula.¹⁴

Suppose $\bar{x} = x_1 \dots x_r$. Given a structure \mathcal{A} and $\bar{a}, \bar{b} \in A^r$, let $d_\psi(\bar{a}, \bar{b})$ be the length of the shortest ψ -path connecting \bar{a} and \bar{b} ,

$$d_\psi(\bar{a}, \bar{b}) := \min\{k > 0 \mid \text{there exist } \bar{a}_0 = \bar{a}, \bar{a}_1, \dots, \bar{a}_k = \bar{b} \text{ such that} \\ \mathcal{A} \models \psi[\bar{a}_i, \bar{a}_{i+1}] \text{ for } i < k\},$$

where $d_\psi(\bar{a}, \bar{b}) := \infty$ in case the set on the right side is empty. Obviously:

(1) If $d_\psi(\bar{a}, \bar{b}) < \infty$ then $0 < d_\psi(\bar{a}, \bar{b}) \leq \|A\|^r$.

Moreover, $\neg[\text{TC}_{\bar{x}, \bar{y}} \psi(\bar{x}, \bar{y})] \bar{s} \bar{t}$ is equivalent to

$$\|\{\bar{v} \mid d_\psi(\bar{s}, \bar{v}) < \infty\}\| = \|\{\bar{v} \mid d_{(\psi(\bar{x}, \bar{y}) \wedge \neg \bar{y} = \bar{t})}(\bar{s}, \bar{v}) < \infty\}\|.$$

We show stepwise that the ingredients of this equality are definable in FO(posTC). Note that the number $\|A\|^r$ has the string $1\tilde{0}$ of length $r+1$ as its $\|A\|$ -adic representation.

(2) For $\varphi(x_1, \dots, x_r, y_1, \dots, y_r) \in \text{FO}(\text{posTC})$ there is an FO(posTC)-formula $\chi(\bar{x}, \bar{y}, z_1, \dots, z_{r+1})$ expressing $d_\varphi(\bar{x}, \bar{y}) \leq \bar{z}$; more precisely: for any ordered \mathcal{A} and $\bar{a}, \bar{b}, \bar{m} \in A$,

$$\mathcal{A} \models \chi[\bar{a}, \bar{b}, \bar{m}] \quad \text{iff} \quad d_\varphi(\bar{a}, \bar{b}) \leq \lceil \bar{m} \rceil.$$

To obtain such a formula χ , we use the TC operator to go through all the tuples of a path leading from \bar{x} to \bar{y} :

$$\chi(\bar{x}, \bar{y}, \bar{z}) := [\text{TC}_{\bar{w} \bar{z}, \bar{w}' \bar{z}'}(\varphi(\bar{w}, \bar{w}') \wedge \bar{z} < \bar{z}')]\bar{x} 0\tilde{0} \bar{y} \bar{z}.$$

(3) For first-order $\varphi(\bar{x}, \bar{y})$ there is an FO(posTC)-formula $\rho_\varphi(\bar{x}, \bar{z})$ expressing $\|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) < \infty\}\| = \bar{z}$.

For a proof we first note that by (1),

$$\|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) < \infty\}\| = \|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) \leq \|A\|^r\}\|.$$

To obtain a formula $\rho_\varphi(\bar{x}, \bar{z})$ we use an inductive definition of the function

$$g(\bar{u}) (= g(\bar{x}, \bar{u})) := \|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) \leq \bar{u}\}\|,$$

$g(1\tilde{0}) (= \|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) \leq \|A\|^r\}\|)$ being the value we are interested in.

By (1), $g(0\tilde{0}) = 0\tilde{0}$. Suppose $g(\bar{u}) = \|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) \leq \bar{u}\}\| = \bar{z}$ where $\bar{u} = u_1 \dots u_{r+1}$. Fix \bar{y} . If $\bar{u} = 0\tilde{0}$ then $d_\varphi(\bar{x}, \bar{y}) \leq \bar{u} + 1$ is equivalent to $\neg\varphi(\bar{x}, \bar{y})$. If $\bar{u} \neq 0\tilde{0}$ then $d_\varphi(\bar{x}, \bar{y}) \leq \bar{u} + 1$ iff there are \bar{z} many \bar{w} such that

¹⁴ The restriction to structures with at least two elements is removed as for FO(LFP) after 8.2.11.

$\bar{w} \neq \bar{y}$, $d_\varphi(\bar{x}, \bar{w}) \leq \bar{u}$, and $\neg\varphi(\bar{w}, \bar{y})$. In the following formula $\chi_\varphi(\bar{x}, \bar{u}, \bar{z}, \bar{z}')$ the inner TC operator serves to count these \bar{w} ($\bar{q} = q_1 \dots q_{r+1}$ being the counting variables) and thus to check whether there are \bar{z} many. If $g(\bar{u}) = \bar{z}$ then $\chi_\varphi(\bar{x}, \bar{u}, \bar{z}, \bar{z}')$ expresses that $g(\bar{u} + 1) = \bar{z}'$. Set with $\bar{v} = v_1 \dots v_{r+1}$

$$\begin{aligned} \chi_\varphi(\bar{x}, \bar{u}, \bar{z}, \bar{z}') &:= [\text{TC}_{\bar{y}\bar{v}\bar{v}'\bar{v}'} ((\bar{y}' = \bar{y} + 1 \wedge \bar{v}' = \bar{v} + 1 \wedge d_\varphi(\bar{x}, \bar{y}) \leq \bar{u} + 1) \\ &\quad \vee (\bar{y}' = \bar{y} + 1 \wedge \bar{v}' = \bar{v} \wedge \bar{u} = \tilde{0} \wedge \neg\varphi(\bar{x}, \bar{y})) \\ &\quad \vee (\bar{y}' = \bar{y} + 1 \wedge \bar{v}' = \bar{v} \wedge \neg\bar{u} = \tilde{0}) \\ &\quad \wedge [\text{TC}_{\bar{w}\bar{q}\bar{w}'\bar{q}'} ((\bar{w}' = \bar{w} + 1 \wedge \bar{q}' = \bar{q}) \vee (\bar{w}' = \bar{w} + 1 \wedge \bar{q}' = \bar{q} + 1 \\ &\quad \wedge \neg\bar{w} = \bar{y} \wedge d_\varphi(\bar{x}, \bar{w}) \leq \bar{u} \wedge \neg\varphi(\bar{w}, \bar{y}))]] \tilde{0} \tilde{0} \tilde{1} \bar{z}']. \end{aligned}$$

By (2), we can view χ_φ as an FO(posTC)-formula. Now

$$\rho_\varphi(\bar{x}, \bar{z}) := [\text{TC}_{\bar{u}\bar{z}\bar{u}'\bar{z}'} (\bar{u}' = \bar{u} + 1 \wedge \chi_\varphi(\bar{x}, \bar{u}, \bar{z}, \bar{z}'))] \tilde{0} \tilde{0} \tilde{1} \bar{z}$$

expresses that $g(\tilde{1}\bar{0}) = \bar{z}$ and hence, that $\|\{\bar{y} \mid d_\varphi(\bar{x}, \bar{y}) < \infty\}\| = \bar{z}$. We can regard ρ_φ as an FO(posTC)-formula. This finishes the proof of (3).

As remarked above, $\neg[\text{TC}_{\bar{x}\bar{y}} \psi(\bar{x}, \bar{y})] \bar{s} \bar{t}$ is equivalent to

$$\|\{\bar{y}' \mid d_\psi(\bar{s}, \bar{y}') < \infty\}\| = \|\{\bar{y}' \mid d_{(\psi(\bar{x}, \bar{y}) \wedge \neg\bar{y} = \bar{t})}(\bar{s}, \bar{y}') < \infty\}\|$$

and thus, by (3), to the FO(posTC)-formula

$$\exists \bar{z} (\rho_\psi(\bar{s}, \bar{z}) \wedge \rho_{(\psi(\bar{x}, \bar{y}) \wedge \neg\bar{y} = \bar{t})}(\bar{s}, \bar{z})). \quad \square$$

In view of 8.6.14 we now have

Corollary 8.6.16 *Let $\tau \supseteq \{<, S, \min, \max\}$. On ordered structures with at least two elements, every FO(TC)-formula is equivalent to a formula of the form*

$$[\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \widetilde{\min \max}$$

where ψ is quantifier-free. □

Exercise 8.6.17 Let τ be arbitrary. Show that for any existential formula φ of FO(posTC) there is a quantifier-free ψ such that in structures with at least two elements, φ and $\exists(\forall)u[\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{u}$ are equivalent. Conclude that, on ordered structures, the same results hold for FO(TC). Hint: Note that in structures with at least two elements any φ of the form

$$(*) \quad [\text{TC}_{\bar{x}, \bar{x}'} \chi(\bar{x}, \bar{x}')] \bar{s} \bar{t}$$

is equivalent to $\exists(\forall)u[\text{TC}_{\bar{x}v\bar{w}, \bar{x}'v'\bar{w}'} \psi(\bar{x}v\bar{w}, \bar{x}'v'\bar{w}')] \tilde{u} \tilde{u}$, where

$$\begin{aligned} \psi(\bar{x}v\bar{w}, \bar{x}'v'\bar{w}') &:= (\bar{x} = \tilde{v} = \tilde{w} \wedge \bar{x}' = \bar{s} \wedge \neg v' = w') \\ &\quad \vee (\neg v = w \wedge \chi(\bar{x}, \bar{x}') \wedge \neg v' = w') \\ &\quad \vee (\bar{x} = \bar{t} \wedge \bar{x}' = \tilde{v} = \tilde{w}). \end{aligned}$$

To show that any existential FO(posTC)-formula φ is equivalent to a formula of the form (*) with quantifier-free χ , get rid of the constants c and d in the proof of 8.6.13: the role of c is taken over by any (all) pairs vw with $v = w$, that of d by any (all) vw with $v \neq w$ (similarly as in ψ). \square

The next example shows that the class of formulas having the normal form of Proposition 8.6.13 does not contain all FO(posTC)-formulas.

Example 8.6.18 Let $\tau = \{E, c, d\}$ where E is binary. There is no sentence of the form

$$(*) \quad [\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{c} \tilde{d},$$

where ψ is first-order, that is equivalent to $\forall x \forall y [\text{TC}_{x,y} Exy] xy$. By contradiction, suppose that (*) is such a sentence. Let $\bar{x} = x_1 \dots x_m$ and set $k := \text{qr}(\psi)$. Choose l big enough compared with k and m . Since $(\mathcal{G}_l, 0, 1) \models \forall x \forall y [\text{TC}_{x,y} Exy] xy$ (for \mathcal{G}_l cf. 2.3.8), we have $(\mathcal{G}_l, 0, 1) \models [\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{c} \tilde{d}$. Hence, there are $\bar{0} = \bar{e}_0, \bar{e}_1, \dots, \bar{e}_n = \bar{1}$ with

$$(\mathcal{G}_l, 0, 1) \models \psi[\bar{e}_i, \bar{e}_{i+1}]$$

for $i < n$. In view of 2.3.8 and the choice of l , we have

$$(\mathcal{G}_l, 0, 1) \dot{\cup} \mathcal{G}_l \models \psi[\bar{e}_i, \bar{e}_{i+1}],$$

where \bar{e}_i, \bar{e}_{i+1} denote the corresponding elements in the first copy. But then

$$(\mathcal{G}_l, 0, 1) \dot{\cup} \mathcal{G}_l \models [\text{TC}_{\bar{x}, \bar{x}'} \psi(\bar{x}, \bar{x}')] \tilde{c} \tilde{d},$$

even though $(\mathcal{G}_l, 0, 1) \dot{\cup} \mathcal{G}_l \not\models \forall x \forall y [\text{TC}_{x,y} Exy] xy$. \square

8.6.3 FO(TC) < FO(LFP)

The preceding example motivates how to handle the TC operation in a game-theoretic characterization of the FO(TC)-equivalence of structures. Given $m, r \geq 1$ and structures \mathcal{A} and \mathcal{B} , the game $\text{TC-}G_m^r(\mathcal{A}, \mathcal{B})$ consists of m moves and is played, as usual, by two players, the spoiler and the duplicator (we shall see below that the number r is related to the arity of the TC operators). In every move there are chosen some elements (possibly more than one) both in \mathcal{A} and \mathcal{B} . Let $a_1, \dots, a_s \in A$ and $b_1, \dots, b_s \in B$ be the elements chosen in the first i moves of a play and denote by p_i the assignment $a_1 \dots a_s \mapsto b_1 \dots b_s$. The duplicator wins the play if p_i is a partial isomorphism for all $i \leq m$ (or, equivalently, if p_m is a partial isomorphism). In the i -th move the spoiler first decides whether the move should be a point move or a TC move; moreover, he chooses one of the structures, say \mathcal{A} . The point moves are as in the first-order case. In a TC move the spoiler selects a j with $1 \leq j \leq r$ and, for some $k \geq 1$, a sequence $\bar{a}_0, \dots, \bar{a}_k$ of j -tuples in \mathcal{A} with

\bar{a}_0 and \bar{a}_k in the domain of p_{i-1} (i.e., \bar{a}_0 and \bar{a}_k consists of elements already chosen or of constants). The duplicator answers by selecting, for some $l \geq 1$, a sequence $\bar{b}_0, \dots, \bar{b}_l$ of j -tuples in \mathcal{B} with $\bar{b}_0 = p_{i-1}(\bar{a}_0)$ and $\bar{b}_l = p_{i-1}(\bar{a}_k)$. Then the spoiler selects some $l' < l$ and the duplicator answers with a $k' < k$. This finishes the i -th move, the elements chosen in this move being $\bar{a}_{k'}, \bar{a}_{k'+1}$ and $\bar{b}_{l'}, \bar{b}_{l'+1}$.

TC moves in which the spoiler chooses the structure \mathcal{A} are called *positive*, those in which he chooses \mathcal{B} , *negative*.

Define the quantifier rank as in the case of FO with the additional clause

$$\text{qr}([\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}) := \text{qr}(\varphi) + 1.$$

Then one can show:

Theorem 8.6.19 *The following are equivalent:*

- (i) *The duplicator has a winning strategy for $\text{TC-G}_m^r(\mathcal{A}, \mathcal{B})$.*
- (ii) *\mathcal{A} and \mathcal{B} satisfy the same $\text{FO}(\text{TC})$ -sentences of quantifier rank $\leq m$ containing the TC operator only in the form $[\text{TC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}$ with $\text{length}(\bar{x}) \leq r$.*

Proof. The proof proceeds along the usual lines; for the implication (ii) \Rightarrow (i) we give a direct proof of the only nontrivial step.

We assume (ii) and explain a winning strategy for the duplicator in $\text{TC-G}_m^r(\mathcal{A}, \mathcal{B})$. Suppose that in the first s moves of a play elements \bar{a} in \mathcal{A} and \bar{b} in \mathcal{B} have been chosen and that

$$(*) \quad \mathcal{A} \models \varphi[\bar{a}] \quad \text{iff} \quad \mathcal{B} \models \varphi[\bar{b}]$$

holds for all $\text{FO}(\text{TC})$ -formulas $\varphi(\bar{x})$ of quantifier rank $\leq m-s$ ($s = 0$ corresponds to the assumption (ii)). Let $p := \bar{a} \mapsto \bar{b}$. Assume that the spoiler decides that the $(s+1)$ -th move should be a TC-move, and let j with $1 \leq j \leq r$ and the sequence $\bar{a}_0, \dots, \bar{a}_k$ of j -tuples be his selections. For $i < k$ set

$$\Psi_i := \{\psi(\bar{x}, \bar{y}, \bar{z}) \mid \text{qr}(\psi) \leq m-s-1, \mathcal{A} \models \psi[\bar{a}, \bar{a}_i, \bar{a}_{i+1}]\}$$

and let $\chi(\bar{x}, \bar{y}, \bar{z}) := \bigvee_{i < k} \bigwedge \Psi_i$. (Of course, $\bigwedge \Psi_i$ is equivalent to a finite conjunction.) Clearly,

$$\mathcal{A} \models [\text{TC}_{\bar{y}, \bar{z}} \chi(\bar{a}, \bar{y}, \bar{z})] \bar{y} \bar{z} [\bar{a}_0, \bar{a}_k].$$

By (*),

$$(*) \quad \mathcal{B} \models [\text{TC}_{\bar{y}, \bar{z}} \chi(\bar{b}, \bar{y}, \bar{z})] \bar{y} \bar{z} [p(\bar{a}_0), p(\bar{a}_k)].$$

The duplicator selects a χ -path $p(\bar{a}_0) = \bar{b}_0, \bar{b}_1, \dots, \bar{b}_l = p(\bar{a}_k)$ witnessing (*). Suppose the spoiler now chooses $\bar{b}_{l'}, \bar{b}_{l'+1}$. Then, for some $k' < k$, $\mathcal{B} \models \Psi_{k'}[\bar{b}, \bar{b}_{l'}, \bar{b}_{l'+1}]$, and the duplicator answers by $\bar{a}_{k'}, \bar{a}_{k'+1}$. By definition of $\Psi_{k'}$,

$$\mathcal{A} \models \psi[\bar{a}, \bar{a}_{k'}, \bar{a}_{k'+1}] \quad \text{iff} \quad \mathcal{B} \models \psi[\bar{b}, \bar{b}_{l'}, \bar{b}_{l'+1}]$$

holds for all $\text{FO}(\text{TC})$ -formulas of quantifier rank $\leq m-s-1$. □

We now aim at a proof of $\text{FO}(\text{TC}) < \text{FO}(\text{LFP})$. The following classes of $\text{FO}(\text{TC})$ -formulas will be useful. For $m \geq 0$ and $r \geq 1$ define P_m^r by induction on m as follows:


- P_0^r is the set of quantifier-free formulas.
- P_{2m+1}^r is the closure under conjunctions and disjunctions of the set $P_{2m}^r \cup \{[\text{TC}_{\bar{x},\bar{y}} \varphi] \bar{s}\bar{t} \mid \varphi \in P_{2m}^r, \text{length}(\bar{x}) \leq r\}$.
- P_{2m+2}^r is the closure under conjunctions and disjunctions of the set $P_{2m+1}^r \cup \{\neg[\text{TC}_{\bar{x},\bar{y}} \neg\varphi] \bar{s}\bar{t} \mid \varphi \in P_{2m+1}^r, \text{length}(\bar{x}) \leq r\}$.

Since $\exists x\varphi(x)$ and $[\text{TC}_{x,y}((x = c \wedge \varphi(y)) \vee (\varphi(x) \wedge y = c))]$ *cc* are equivalent, one gets by a simple induction on formulas:

Lemma 8.6.20 *If the vocabulary τ contains at least one constant then every formula of $\text{FO}(\text{TC})$ is equivalent to a formula in $\bigcup_{m \geq 0, r \geq 1} P_m^r$. \square*

We show $\text{FO}(\text{TC}) < \text{FO}(\text{LFP})$ by exhibiting a class of structures axiomatizable in $\text{FO}(\text{LFP})$ but not in $\text{FO}(\text{TC})$. The underlying idea is already present in the following example.

Example 8.6.21 Let $r \geq 1$ and $\tau = \{E, F, c\}$ with E binary and F (“full”) unary. Let \mathcal{A} and \mathcal{B} be the structures



Points for which F holds are represented by full dots, $\overset{a}{\downarrow} b$ means that Eab . We show for quantifier-free φ and \bar{x} with $\text{length}(\bar{x}) \leq r$ that

$$\mathcal{A} \models [\text{TC}_{\bar{x},\bar{y}} \varphi] \tilde{c}\tilde{c} \quad \text{implies} \quad \mathcal{B} \models [\text{TC}_{\bar{x},\bar{y}} \varphi] \tilde{c}\tilde{c}.$$

In fact, suppose that $\tilde{c} = \bar{a}_0, \dots, \bar{a}_l = \tilde{c}$ witnesses $(\tilde{c}, \tilde{c}) \in \text{TC}(\varphi^{\mathcal{A}}(-, -))$. Successively choose $\bar{b}_0 = \tilde{c}, \bar{b}_1, \dots, \bar{b}_l = \tilde{c}$ in \mathcal{B} such that $\bar{a}_i \bar{a}_{i+1} \mapsto \bar{b}_i \bar{b}_{i+1}$ is a partial isomorphism (such elements exist, since \mathcal{B} contains $2r$ full elements). As φ is quantifier-free, $\mathcal{B} \models \varphi[\bar{b}_i, \bar{b}_{i+1}]$ and hence, $(\tilde{c}, \tilde{c}) \in \text{TC}(\varphi^{\mathcal{B}}(-, -))$. \square


Using suitable refinements of \mathcal{A}, \mathcal{B} one can handle nested TC operators and show:

Theorem 8.6.22 $\text{FO}(\text{TC}) < \text{FO}(\text{LFP})$.


Proof. Fix an $r \geq 1$ and let $\tau_0 := \{E, F\}$, where E and F are as in the preceding example. For $m \geq 0$ define τ_0 -structures $\mathcal{A}_{r,m}$ and $\mathcal{B}_{r,m}$ and their roots by induction on m as follows:

$$\mathcal{A}_{r,0} : \quad \circ \qquad \mathcal{B}_{r,0} : \quad \bullet$$

and the respective element is the root. Suppose that $\mathcal{A}_{r,m}$ and $\mathcal{B}_{r,m}$ and their roots have already been defined. Set



that is, $\mathcal{A}_{r,m+1}$ is the disjoint union of $r + 1$ copies of $\mathcal{B}_{r,m}$ together with a new element, the root of $\mathcal{A}_{r,m+1}$, that is connected by the relation E to the roots of the distinct copies of $\mathcal{B}_{r,m}$. The structure $\mathcal{B}_{r,m+1}$ is defined in the same way, but one copy of $\mathcal{B}_{r,m}$ is replaced by $\mathcal{A}_{r,m}$. Note that only leaves can be full. The structure $\mathcal{B}_{2,2}$ looks like



Let $\tau = \tau_0 \cup \{c\}$ and let $\mathcal{A}'_{r,m}$ and $\mathcal{B}'_{r,m}$ be the τ -structures ($\mathcal{A}_{r,m}$, “root of $\mathcal{A}_{r,m}$ ”) and ($\mathcal{B}_{r,m}$, “root of $\mathcal{B}_{r,m}$ ”), respectively.

We recall or introduce some terminology. Let \mathcal{D} be one of the structures $\mathcal{A}_{r,m}$ or $\mathcal{B}_{r,m}$. For $d, d' \in \mathcal{D}$ we say that d' is a *successor* of d if $E^{\mathcal{D}}dd'$, and that d' is a *descendant* of d if for some $l \geq 0$ there are $d_0 = d, d_1, \dots, d_l = d'$ such that $E^{\mathcal{D}}d_i d_{i+1}$ for $i < l$. The *height* $h(d)$ of d is the length of the largest E -path starting at d , i.e.,

$$h(d) := \max\{l \geq 0 \mid \text{there are } d_0 = d, d_1, \dots, d_l \text{ with } E^{\mathcal{D}}d_i d_{i+1} \text{ for } i < l\}.$$

The substructure $\mathcal{T}(d)$ of \mathcal{D} whose universe $T(d)$ is the set of descendants of d is called the *tree generated by* d .


Define the function colour: $D \rightarrow \{\text{black, white}\}$ by

$$(*) \quad \text{colour}(d) := \begin{cases} \text{black} & \text{if } \mathcal{T}(d) \cong \mathcal{B}_{r,h(d)} \\ \text{white} & \text{otherwise.} \end{cases}$$

The inductive definition of $\mathcal{A}_{r,m}$ and $\mathcal{B}_{r,m}$ leads to the following definition of colour by induction on the height: The point d is of black colour if it is full (i.e., if $F^D d$ holds) or if it has a white successor. More precisely, $\text{colour}(d) = \text{black}$ iff

$$D \models [\text{IFP}_{x,X}(Fx \vee \exists y \exists z (Exy \wedge Exz \wedge Xy \wedge \neg Xz))] x [d].$$

Using \ominus for a black point and \circ for a white point, the structure $\mathcal{B}_{2,2}$ looks as follows:



Let φ_0 be a sentence of $\text{FO}(\text{LFP})[\tau]$ saying that the root is black, that is, an $\text{FO}(\text{LFP})$ -sentence equivalent to $[\text{IFP}_{x,X}(Fx \vee \exists y \exists z (Exy \wedge Exz \wedge Xy \wedge \neg Xz))] c$. Then, in view of $(*)$, we have

$$(1) \text{ For all } m \geq 0, r \geq 1 : \mathcal{A}'_{r,m} \models \neg \varphi_0 \text{ and } \mathcal{B}'_{r,m} \models \varphi_0.$$

Therefore, by Lemma 8.6.20, we are done if we show:

$$(2) \text{ Let } k \geq 0 \text{ and let } \varphi \text{ be an } \text{FO}(\text{TC})\text{-sentence in } P_k^r. \text{ Then for } m > 2 \cdot k,$$

$$\mathcal{A}'_{2r,m} \models \varphi \quad \text{iff} \quad \mathcal{B}'_{2r,m} \models \varphi.$$

We prove (2) by induction on k , showing a corresponding statement for formulas and assignments to their variables. As in the example preceding this theorem one carefully has to choose the elements in the right copies. To make precise the notion of “right copy” we have in mind, we introduce the concept of an *amenable* function. A function $\pi : \mathcal{A}_{r,m} \rightarrow \mathcal{B}_{r,m}$ is *amenable* if

- $\pi : \mathcal{A}_{r,m} \upharpoonright \{E\} \cong \mathcal{B}_{r,m} \upharpoonright \{E\}$
- For every $l \leq m$ there exists exactly one $a \in \mathcal{A}_{r,m}$ such that $h(a) = l$ and $\text{colour}(\pi(a)) \neq \text{colour}(a)$. The element a is called the π -critical point of height l .

Note that if π is amenable, a its critical point of height l , a' its critical point of height l' , and $l' < l$, then $a' \in T(a)$. There is an amenable function from

$\mathcal{A}_{r,m}$ to $\mathcal{B}_{r,m}$; this is obtained by induction on m : If $\pi_0 : \mathcal{A}_{r,m} \rightarrow \mathcal{B}_{r,m}$ is amenable, then $\pi : \mathcal{A}_{r,m+1} \rightarrow \mathcal{B}_{r,m+1}$ is amenable, where (compare the figures defining $\mathcal{A}_{r,m+1}$ and $\mathcal{B}_{r,m+1}$) π maps the root of $\mathcal{A}_{r,m+1}$ to the root of $\mathcal{B}_{r,m+1}$, π equals π_0^{-1} on the leftmost copy of $\mathcal{B}_{r,m}$ that is attached to the root of $\mathcal{A}_{r,m+1}$, and is the identity on the remaining copies.

In the example preceding this theorem, when choosing the elements \bar{b}_i , we avoided the element of \mathcal{B} , which is not full and distinct from the root. Here we have to avoid elements below critical points of suitable height. More precisely, one can show by induction on k :

(3) Suppose that $s, k \geq 0$ and $m > 2 \cdot k$. Let $\bar{a}_1, \bar{a}'_1, \dots, \bar{a}_s, \bar{a}'_s$ be r -tuples in $\mathcal{A}_{2r,m}$. Furthermore, let $\pi : \mathcal{A}_{2r,m} \rightarrow \mathcal{B}_{2r,m}$ be amenable and let a be its critical point of height $2k$. Finally, assume that

$$\{\bar{a}_1, \bar{a}'_1, \dots, \bar{a}_s, \bar{a}'_s\} \cap T(a) = \emptyset.$$

Then, for every $\varphi(\bar{x}_1, \bar{x}'_1, \dots, \bar{x}_s, \bar{x}'_s) \in P_k^r$ with free variables among the displayed ones,

$$\mathcal{A}'_{2r,m} \models \varphi[\bar{a}_1, \bar{a}'_1, \dots, \bar{a}_s, \bar{a}'_s] \text{ iff } \mathcal{B}'_{2r,m} \models \varphi[\pi(\bar{a}_1), \pi(\bar{a}'_1), \dots, \pi(\bar{a}_s), \pi(\bar{a}'_s)].$$

Once this is shown, we get (2) by taking $s := 0$ and as π any amenable function from $\mathcal{A}_{2r,m}$ to $\mathcal{B}_{2r,m}$.

In the inductive proof of (3) the case $k = 0$ is obvious. To prove (3) for $k \geq 1$ under the assumption that it is already proved for $k - 1$, it suffices to show for $m > 2 \cdot k$ (writing \bar{a} for $\bar{a}_1 \bar{a}'_1 \dots \bar{a}_s \bar{a}'_s$)

$$(+) \quad \mathcal{A}'_{2r,m} \models [\text{TC}_{\bar{x}, \bar{y}} \psi] \bar{s} \bar{t} [\bar{a}] \text{ iff } \mathcal{B}'_{2r,m} \models [\text{TC}_{\bar{x}, \bar{y}} \psi] \bar{s} \bar{t} [\pi(\bar{a})]$$

where $\text{length}(\bar{x}) = r$ and $\psi \in P_{k-1}^r$.

First we assume that the left side of (+) holds, $\bar{c}_0, \dots, \bar{c}_l$ being a witnessing sequence. Our aim is to construct a witnessing sequence for the right side. As \bar{c}_1 contains less than $2r+1$ elements, not all of the $2r+1$ copies of $\mathcal{B}_{2r,2k-1}$ attached to a contain one of these elements. Hence, we can find \bar{d}_1 in $\mathcal{B}_{2r,m}$, but outside the copy of $\mathcal{A}_{2r,2k-1}$ attached to $\pi(a)$, and an amenable $\pi' : \mathcal{A}'_{2r,m} \rightarrow \mathcal{B}'_{2r,m}$ that coincides with π outside of $T(a) \setminus \{a\}$ such that

$$\pi'(\bar{c}_1) = \bar{d}_1$$

and such that $\bar{a}, \bar{c}_0, \bar{c}_1$ and π' satisfy (3) for $k - 1$. Hence, by induction hypothesis,

$$\mathcal{B}'_{2r,m} \models \psi[\pi'(\bar{a}), \pi'(\bar{c}_0), \pi'(\bar{c}_1)].$$

Treating now the pair \bar{c}_1, \bar{c}_2 in the same way (\bar{c}_1, \bar{c}_2 contains less than $2r+1$ elements!), where the counterpart \bar{d}_1 of \bar{c}_1 is already defined, and continuing in this way up to the pair \bar{c}_{l-1}, \bar{c}_l , we get a witnessing sequence $\bar{d}_0, \bar{d}_1, \dots, \bar{d}_l$ for the right side of (+).

For the other direction we argue similarly. Concerning elements of the witnessing sequence in the copy of $\mathcal{A}_{2r,2k-1}$ attached to $\pi(a)$, we proceed with this copy as with $T(a)$ above, thus, altogether, decreasing the height of the critical point of the induction hypothesis by two. \square

Notes 8.6.23 Theorem 8.6.9 is due to Grädel and McColm [56]. Subsection 8.6.2 is based on [93, 94]. The trees in subsection 8.6.3 go back to [18].

8.7 Bounded Fixed-Point Logic

In the preceding section we already mentioned that an application of the TC operator can be replaced by an application of the LFP operator. In fact,

$$\models_{\text{fin}} [\text{TC}_{\bar{x},\bar{y}}\varphi(\bar{x},\bar{y})] \bar{s}\bar{t} \leftrightarrow [\text{LFP}_{\bar{y},Y}(\varphi(\bar{s},\bar{y}) \vee \exists \bar{x}(Y\bar{x} \wedge \varphi(\bar{x},\bar{y})))] \bar{t}.$$

Conversely, every FO(LFP)-formula of the form

$$[\text{LFP}_{\bar{y},Y}(\psi(\bar{y}) \vee \exists \bar{x}(Y\bar{x} \wedge \varphi(\bar{x},\bar{y})))] \bar{t}$$

where Y does not occur in ψ and φ , is equivalent to

$$\psi(\bar{t}) \vee \exists \bar{x}(\psi(\bar{x}) \wedge [\text{TC}_{\bar{x},\bar{y}}\varphi(\bar{x},\bar{y})] \bar{x}\bar{t}).$$

In this sense, transitive closures correspond to those least fixed-points, where each tuple in a new stage is already witnessed by a *single* tuple of the preceding stage. We generalize transitive closure logic to *Bounded Fixed-Point Logic* FO(BFP), which allows the LFP operator only if there is a bound $r \geq 1$ such that each tuple in a new stage is already witnessed by a set of at most r many tuples of the preceding stage. More precisely: For a vocabulary τ the class FO(BFP)[τ] of formulas of FO(BFP) of vocabulary τ is given by

- the rules of first-order logic starting from first-order atomic formulas (i.e., no relation variables are allowed)
- the “BFP-rule”

$$\frac{\varphi_0(\bar{y}), \varphi_1(\bar{y}, \bar{x}_1, \dots, \bar{x}_r)}{[\text{LFP}_{\bar{y},Y}(\varphi_0(\bar{y}) \vee \exists \bar{x}_1 \dots \exists \bar{x}_r (Y\bar{x}_1 \wedge \dots \wedge Y\bar{x}_r \wedge \varphi_1(\bar{y}, \bar{x}_1, \dots, \bar{x}_r)))] \bar{t}}$$

where Y is an s -ary relation variable for some $s \geq 1$, and where $r \geq 0$ and $\bar{x}_1, \dots, \bar{x}_r$ are s -tuples of variables, all variables in $\bar{x}_1, \dots, \bar{x}_r$ being pairwise distinct. (Note that Y does not occur free in φ_0 or φ_1 .)

We abbreviate the formula above by

$$[\text{LFP}_{\bar{y},Y}(\varphi_0(\bar{y}) \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_r \in Y \varphi_1(\bar{y}, \bar{x}_1, \dots, \bar{x}_r))] \bar{t}.$$

Sometimes, bounded fixed-point logic is called *stratified fixed-point logic*. This name will become clear in the next chapter.

Denote by $\text{FO}(\text{BFP}_{r_0})$ the logic obtained by restricting the BFP-rule to r with $r \leq r_0$. The introductory remarks show part (c) of the following proposition, parts (a) and (b) being trivial.

- Proposition 8.7.1** (a) $\text{FO}(\text{BFP}) \leq \text{FO}(\text{LFP})$
 (b) $\text{FO}(\text{BFP}_0) \equiv \text{FO}$
 (c) $\text{FO}(\text{BFP}_1) \equiv \text{FO}(\text{TC})$
 (d) $\text{FO}(\text{BFP}_2) \equiv \text{FO}(\text{BFP})$.

Proof. (d) The proof proceeds by induction on $\text{FO}(\text{BFP})$ -formulas. In the main step we show that

$$(*) \quad [\text{LFP}_{\bar{y}, Y} (\varphi_0(\bar{y}) \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_r \in Y \varphi_1(\bar{y}, \bar{x}_1, \dots, \bar{x}_r))] \bar{t}$$

is equivalent to a formula of $\text{FO}(\text{BFP}_2)$ under the assumption that φ_0 and φ_1 are already in $\text{FO}(\text{BFP}_2)$. For simplicity we assume that Y is unary and hence, $\bar{y} = y, \bar{x}_1 = x_1, \dots, \bar{x}_r = x_r$. Denote by Y_0, Y_1, \dots the stages corresponding to the LFP operation in $(*)$. Let Z be a new r -ary relation symbol. The fixed-point operation of the following formula ψ stepwise incorporates the elements of a new stage Y_n into the different components of the tuples in Z .

$$\psi(y_1, \dots, y_r) := [\text{LFP}_{\bar{z}, Z} ((\varphi_0(z_1) \wedge \dots \wedge \varphi_0(z_r)) \vee \exists \bar{x} \in Z \exists \bar{y} \in Z \exists y (\varphi_1(y, \bar{x}) \wedge \bigvee_{i=1, \dots, r} (z_i = y \wedge \bigwedge_{j \neq i} z_j = y_j))] \bar{y}.$$

If Z_0, Z_1, \dots are the stages of the LFP operation in ψ , then one easily shows by induction on n :

- (1) for all n , $Z_n \subseteq Y_n^r$
 (2) for all n there is an m such that $Y_n^r \subseteq Z_m$.

Hence, $Y_\infty^r = Z_\infty$, and therefore,

$$\models_{\text{fin}} [\text{LFP}_{y, Y} (\varphi_0 \vee \exists x_1 \in Y \dots \exists x_r \in Y \varphi_1)] t \leftrightarrow \psi(t, \dots, t). \quad \square$$

We come back to the trees $\mathcal{A}_{r,m}$ and $\mathcal{B}_{r,m}$ introduced in the proof of 8.6.22. A point of one of these structures got the colour “black”, if it was (a leaf and) full or if at least (= exactly) one successor was white. Thus the colour of a point was determined by the colour of its $r + 1$ successors. It is therefore not surprising that we can find a sentence in $\text{FO}(\text{BFP}_{r+1})$ which separates the classes $\{\mathcal{A}'_{r,m} \mid m \geq 0\}$ and $\{\mathcal{B}'_{r,m} \mid m \geq 0\}$. In fact, we have (recall that c denotes the root of the structures in question)

$$F_\infty^{\psi_r} = \{(c, y) \mid \text{colour}(y) = \text{black}\} \cup \{(x, y) \mid x \neq c, \text{colour}(y) = \text{white}\}$$

for the formula $\psi_r(x, y, Y) :=$

$$\begin{aligned} & ((x = c \wedge Fy) \vee (\neg x = c \wedge \forall z \neg Eyz \wedge \neg Fy)) \vee \exists x_0 y_0 \in Y \dots \exists x_r y_r \in Y \\ & \left(\bigwedge_{0 \leq i < j \leq r} y_i \neq y_j \wedge \bigwedge_{i \leq r} Eyy_i \wedge \right. \\ & \left. ((x = c \wedge \bigvee_{i \leq r} \neg x_i = c) \vee (\neg x = c \wedge \bigwedge_{i \leq r} x_i = c)) \right). \end{aligned}$$

Hence for all $m > 0$,

$$\mathcal{A}'_{r,m} \models \neg[\text{LFP}_{xy,Y}\psi_r]cc \text{ and } \mathcal{B}'_{r,m} \models [\text{LFP}_{xy,Y}\psi_r]cc.$$

It can be shown that (2) in the proof of 8.6.22 can be strengthened to

for any FO(TC)-sentence φ there is an m_0 such that for all $m \geq m_0$,

$$\mathcal{A}'_{2,m} \models \varphi \text{ iff } \mathcal{B}'_{2,m} \models \varphi.$$

Therefore, the FO(BFP)-sentence $[\text{LFP}_{xy,Y}\psi_2]cc$ is not equivalent to an FO(TC)-sentence. This gives the first inequality in the next proposition. To obtain the second inequality one shows that the FO(LFP)-sentence φ_0 of (1) in the proof of 8.6.22 is not equivalent to an FO(BFP)-sentence; that is, there is no FO(BFP)-sentence separating the sets $\{\mathcal{A}'_{r,m} \mid r \geq 1, m \geq 0\}$ and $\{\mathcal{B}'_{r,m} \mid r \geq 1, m \geq 0\}$.

Theorem 8.7.2 FO(TC) < FO(BFP) < FO(LFP). □

We shall see in Theorem 9.1.6 that FO(BFP) \equiv FO(LFP) on ordered structures. In the BFP-rule the fixed-point operation is applied to formulas $(\varphi_0 \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_r \in Y \varphi_1)$ which do not contain free relation variables besides Y , and where Y only occurs positively. In fact, this is the only essential restriction, as is shown by the following result that we need in the next chapter.

Proposition 8.7.3 *Let Y be a second-order variable and let ψ be a formula built up from FO(BFP $_r$)-formulas and atomic formulas of the form $Y\bar{s}$ with \neg, \wedge, \vee , and \exists . Furthermore, suppose that Y is not in the scope of any negation sign. If Y occurs at most r times in ψ , then $[\text{LFP}_{\bar{y},Y}\psi]\bar{t}$ is equivalent to a formula of FO(BFP $_r$).*

Proof. To each such ψ with exactly s occurrences of Y we associate a formula ψ^* of the form $\exists \bar{x}_1 \in Y \dots \exists \bar{x}_n \in Y \psi_1$, where ψ_1 is an FO(BFP $_r$)-formula that does not contain Y and is equivalent to ψ whenever a nonempty relation is assigned to Y . Then, $[\text{LFP}_{\bar{y},Y}\psi(\bar{y}, Y)]\bar{t}$ is equivalent to

$$[\text{LFP}_{\bar{y},Y}(\psi(\bar{y}, \emptyset) \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_n \in Y \psi_1)]\bar{t},$$

a formula of FO(BFP $_r$). We define ψ^* by induction on ψ : If $\psi = Y\bar{s}$ then $\psi^* := \exists \bar{x} \in Y \bar{x} = \bar{s}$. Set $\psi^* := \psi$ for any formula not containing Y ; in particular, for $\psi = \neg\varphi$ we have $\psi^* = \neg\varphi$ since, by hypothesis, φ does not contain

Y . The definition of ψ^* for the other cases is straightforward; it uses that $(\varphi \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_n \in Y \psi)$ and $\exists \bar{x}_1 \in Y \dots \exists \bar{x}_n \in Y (\varphi \vee \psi)$ are equivalent in case no variable in $\bar{x}_1, \dots, \bar{x}_n$ is free in φ and provided a nonempty relation is assigned to Y (and a similar equivalence for the conjunction). \square

Notes 8.7.4 Stratified logic was introduced and studied first by Chandra and Harel [19]. References for the results given are [23, 102, 61].

9. Logic Programs

In the present chapter we introduce a group of languages that originate in database theory. Our aim is to exhibit a close relationship to fixed-point logics, thus demonstrating that the two groups of logical systems are two sides of one and the same coin. We therefore may use methods and results from one group when investigating the other. Examples will illustrate the methodological means we thus have gained.

9.1 DATALOG

Fix a vocabulary τ . A *general logic program* is a finite set Π of *clauses* of the form

$$\gamma \leftarrow \gamma_1, \dots, \gamma_l$$

where $l \geq 0$ and where $\gamma, \gamma_1, \dots, \gamma_l$ are atomic or negated atomic first-order formulas. The formula γ must be an atomic formula of the form $R\bar{t}$. It is called the *head* of the clause. The sequence $\gamma_1, \dots, \gamma_l$ constitutes the *body* of the clause. Every relation symbol occurring in the head of some clause of Π is *intentional*, all the other symbols in τ are *extensional*. Denote by $(\tau, \Pi)_{\text{int}}$ and $(\tau, \Pi)_{\text{ext}}$ the set of intentional and extensional symbols, respectively. Hence, $(\tau, \Pi)_{\text{ext}} = \tau \setminus (\tau, \Pi)_{\text{int}}$. We often write τ_{int} and τ_{ext} if no confusion is possible.

A *DATALOG program* is a general logic program, in which no intentional symbol occurs negated in the body of any clause. Examples of DATALOG programs are

$$\begin{array}{lll} \Pi_0 : & Txy \leftarrow Exy & \Pi_1 : & P \text{ min} \leftarrow \\ & Txx \leftarrow Txy, Eyz & & Qy \leftarrow Px, Sxy \\ & & & Py \leftarrow Qx, Sxy \end{array} \quad \Pi_2 : \quad \begin{array}{l} Px \leftarrow \neg x = c \\ Px \leftarrow \neg Rx \end{array}$$

The intended meaning of a DATALOG program Π is obtained by interpreting its clauses in a dynamic way as rules, the clause $\gamma \leftarrow \gamma_1, \dots, \gamma_l$ as the rule

$$(*) \quad \text{whenever } \gamma_1, \dots, \gamma_l \text{ then } \gamma.$$

The values of the extensional symbols are given by a τ_{ext} -structure. All the rules of Π will be applied simultaneously to generate consecutive stages of

the intentional symbols. For Π_0 for example, the final value of T in a graph $\mathcal{G} = (G, E^G)$ will be the transitive closure of E^G . We come to the exact definition. To simplify the presentation of the semantics we suppose that for any intentional symbol P of Π , say, of arity r , there are distinct variables $\bar{x}_P = x_{P,1} \dots x_{P,r}$ such that any clause in Π with head symbol P has the form

$$P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l$$

(otherwise, carry out suitable changes having in mind the intended interpretation; for example, replace the clause $P \min \leftarrow$ in Π_1 by $Px \leftarrow x = \min$ (where $\bar{x}_P = x$), and a clause $Rxyx \leftarrow Py, Sx \min$ by $Rxyz \leftarrow Py, Sx \min, x = z$ (where $\bar{x}_R = xyz$)). With every intentional P we associate the formula

$$\varphi_P(\bar{x}_P) := \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \}$$

where \bar{y} contains the variables in $\gamma_1, \dots, \gamma_l$ that are distinct from the variables in \bar{x}_P . Assume that P^1, \dots, P^k are the intentional symbols of Π . By definition of a DATALOG program they do not occur negated in the body of any clause. Therefore $\varphi_{P^1}, \dots, \varphi_{P^k}$ are positive in P^1, \dots, P^k and hence, their simultaneous fixed-point $(P^1_{(\infty)}, \dots, P^k_{(\infty)})$ given by $P^i_{(\infty)} := \bigcup_{n \geq 0} P^i_{(n)}$ exists,

where

$$P^i_{(0)} := \emptyset, \quad P^i_{(n+1)} := \{ \bar{x}_{P^i} \mid \varphi_{P^i}(\bar{x}_{P^i}, P^1_{(n)}, \dots, P^k_{(n)}) \}.$$

For a τ_{ext} -structure \mathcal{A} we can rewrite the definition of $P^i_{(n)}$ in such a way that the procedural character of the DATALOG program Π is emphasized:

$$P^i_{(0)} = \emptyset$$

$$(+) \quad P^i_{(n+1)} = \{ \bar{a} \mid \text{there are } P^i \bar{x}_{P^i} \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \text{ and } \bar{b} \in A \text{ such that } (\mathcal{A}, P^1_{(n)}, \dots, P^k_{(n)}) \models \gamma_1 \wedge \dots \wedge \gamma_l [\bar{a}, \bar{b}] \}$$

(where \bar{a} interprets the variables in \bar{x}_{P^i} and \bar{b} the remaining ones).

Thus, a DATALOG program Π and a τ_{ext} -structure \mathcal{A} give rise to the τ -structure

$$\mathcal{A}[\Pi] := (\mathcal{A}, P^1_{(\infty)}, \dots, P^k_{(\infty)}).$$

And every intentional symbol P of Π defines a query on the class of τ_{ext} -structures given by $\mathcal{A} \xrightarrow{\Pi} P\mathcal{A}[\Pi]$.

For the DATALOG program Π_1 at the beginning of this section and any ordering \mathcal{B} as $\{<, S, \min, \max\}$ -structure, we have $\mathcal{B}[\Pi_1] = (\mathcal{B}, P^{\mathcal{B}}, Q^{\mathcal{B}})$, where $P^{\mathcal{B}}$ and $Q^{\mathcal{B}}$ consist of the elements at odd and even positions, respectively.

A DATALOG *formula* has the form $(\Pi, P)\bar{t}$ where P is an intentional symbol of Π , say of arity r , and $\bar{t} = t_1 \dots t_r$ are terms. It is a formula of vocabulary

τ_{ext} , its free variables being those occurring in some t_i . For a τ_{ext} -structure \mathcal{A} and $\bar{a} \in A$,

$$\mathcal{A} \models (\Pi, P)\bar{t}[\bar{a}] \quad \text{iff} \quad (t_1[\bar{a}], \dots, t_r[\bar{a}]) \in P^{\mathcal{A}[\Pi]}.$$

In order to get sentences and thus to be able to compare the expressive power of DATALOG programs with other logics, it is desirable to allow zero-ary relation symbols. If $P \in \tau$ is 0-ary then, in a structure \mathcal{A} , $P^{\mathcal{A}}$ is the boolean value TRUE or FALSE. Denoting the empty sequence by \emptyset , for zero-ary P^i the inductive step (+) above reads as follows:

$$P^i_{(n+1)} := \{ \emptyset \mid \text{there are a clause } P^i \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \text{ and } \bar{b} \in A \\ \text{such that } (\mathcal{A}, P^1_{(n)}, \dots, P^k_{(n)}) \models \gamma_1 \wedge \dots \wedge \gamma_l [\emptyset, \bar{b}] \}.$$

So, “ $P^{\mathcal{A}} = \emptyset$ ” means “ $P^{\mathcal{A}} = \text{FALSE}$ ”, “ $P^{\mathcal{A}} = \{ \emptyset \}$ ” means “ $P^{\mathcal{A}} = \text{TRUE}$ ”.

To give an example, consider the program

$$\begin{aligned} \Pi_4 : \quad & Txy \leftarrow Exy \\ & Txz \leftarrow Txy, Tyz \\ & P \leftarrow Tcd \end{aligned}$$

with a 0-ary relation symbol P . Then, for any graph $\mathcal{G} := (G, E^G, c^G, d^G)$ with distinguished points c^G and d^G , $P^{\mathcal{G}[\Pi_4]}$ is true just in case c^G and d^G are connected by a path; in other words, $\mathcal{G} \models (\Pi_4, P)$ iff c^G and d^G are connected by a path.

We now come to a generalization of DATALOG programs. Let Π be a DATALOG program and \mathcal{A} a $(\tau, \Pi)_{\text{ext}}$ -structure. Then $\mathcal{A}[\Pi]$ is a τ -structure. To $\mathcal{A}[\Pi]$ one can apply a further DATALOG program. In such a program all symbols in τ are extensional and thus may occur negated in the bodies of clauses. An iteration of this process leads to the notion of a stratified DATALOG program. A *stratified* DATALOG program (by short: S-DATALOG program) Σ consists of a sequence Π_0, \dots, Π_n of DATALOG programs of vocabularies τ_0, \dots, τ_n , where $(\tau_{m+1}, \Pi_{m+1})_{\text{ext}} = \tau_m$ for $m < n$. Given a $(\tau_0, \Pi_0)_{\text{ext}}$ -structure \mathcal{A} , we set

$$\mathcal{A}[\Sigma] := (\dots(\mathcal{A}[\Pi_0])[\Pi_1] \dots)[\Pi_n].$$

As with DATALOG, $(\Sigma, P)\bar{t}$ is an S-DATALOG *formula* if P is an intentional symbol of one of the constituents Π_i . For example, $\Sigma := \Pi_0, \Pi_1$ with

$$\Pi_0 : \quad Px \leftarrow Exy \quad \Pi_1 : \quad Lx \leftarrow \neg Px.$$

is an S-DATALOG program. In any tree, the S-DATALOG formula $(\Sigma, L)x$ defines the set of leaves.

The following result explains why bounded fixed-point logic is sometimes called stratified fixed-point logic.

Theorem 9.1.1 S-DATALOG \equiv FO(BFP).

We give the proof in such a way that it yields a stronger result: The *breadth* of a DATALOG program is the number of occurrences of intentional symbols in the bodies. Thus the program

$$\begin{aligned} Txy &\leftarrow Exy \\ Txz &\leftarrow Txy, Tyz \\ P &\leftarrow Tcd \end{aligned}$$

has breadth 3. The *breadth* of an S-DATALOG program $\Pi_0, \Pi_1, \dots, \Pi_k$ is, by definition, given by $\max_{i \leq k} \text{breadth}(\Pi_i)$. Denote by S-DATALOG $_r$ the class of S-DATALOG formulas $(\Sigma, Q)\bar{t}$ where Σ is of breadth $\leq r$. Then the preceding theorem can be strengthened to:

Lemma 9.1.2 For all $r \geq 0$, S-DATALOG $_r \equiv$ FO(BFP $_r$).

Proof. We show FO(BFP $_r$) \leq S-DATALOG $_r$ by assigning to every formula φ of FO(BFP $_r$) with free variables among \bar{y} an equivalent S-DATALOG $_r$ formula $(\Sigma_\varphi, Q_\varphi)\bar{y}$. We do this by induction on φ , tacitly assuming that the intentional symbols of the programs $\Pi_\psi, \Pi_\chi, \dots$ we refer to are distinct.

- If $\varphi(\bar{y})$ is atomic, set $\Sigma_\varphi := \{Q_\varphi\bar{y} \leftarrow \varphi\}$.
- If $\varphi(\bar{y}) = \neg\psi$, set $\Sigma_\varphi := \Sigma_\psi, \{Q_\varphi\bar{y} \leftarrow \neg Q_\psi\bar{y}\}$.
- If $\varphi(\bar{y}) = (\psi \vee \chi)$, set $\Sigma_\varphi := \Sigma_\psi, \Sigma_\chi, \{Q_\varphi\bar{y} \leftarrow Q_\psi\bar{y}, Q_\varphi\bar{y} \leftarrow Q_\chi\bar{y}\}$.
- If $\varphi(\bar{y}) = \exists x\psi(\bar{y}, x)$, set $\Sigma_\varphi := \Sigma_\psi, \{Q_\varphi\bar{y} \leftarrow Q_\psi\bar{y}x\}$.
- For $\varphi(\bar{y}) = [\text{LFP}_{\bar{x}, X}(\psi(\bar{y}, \bar{x}) \vee \exists \bar{x}_1 \in X \dots \exists \bar{x}_s \in X \chi(\bar{y}, \bar{x}, \bar{x}_1, \dots, \bar{x}_s))] \bar{t}$ with $s \leq r$ let

$$\begin{aligned} \Pi : \quad Q\bar{y}\bar{x} &\leftarrow Q_\psi\bar{y}\bar{x} \\ Q\bar{y}\bar{x} &\leftarrow Q\bar{y}\bar{x}_1, \dots, Q\bar{y}\bar{x}_s, Q_\chi\bar{y}\bar{x}\bar{x}_1 \dots \bar{x}_s \end{aligned}$$

and set $\Sigma_\varphi := \Sigma_\psi, \Sigma_\chi, \Pi, \{Q_\varphi\bar{y} \leftarrow Q\bar{y}\bar{t}\}$.

For the direction S-DATALOG $_r \leq$ FO(BFP $_r$) we first study the case of a (plain) DATALOG formula $(\Pi, P)\bar{t}$ where Π is of breadth $\leq r$. Let P^1, \dots, P^k with $P = P^1$ be its intentional symbols. By previous remarks we may assume that the clauses have the form

$$Q\bar{x}_Q \leftarrow \gamma_1, \dots, \gamma_l$$

where \bar{x}_Q is a sequence of distinct variables. Let

$$\varphi_Q(\bar{x}_Q) := \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid Q\bar{x}_Q \leftarrow \gamma_1 \dots, \gamma_l \text{ in } \Pi \}$$

where \bar{y} contains all the variables in the clause that are not in \bar{x}_Q . By definition of the semantics of DATALOG, the formula $(\Pi, P)\bar{t}$ is equivalent to

$$[\text{S-LFP}_{\bar{x}_{P^1, P^1}, \dots, \bar{x}_{P^k, P^k}} \varphi_{P^1}, \dots, \varphi_{P^k}] \bar{t}.$$

Since $\varphi_{P^1}, \dots, \varphi_{P^k}$ are existential and altogether, P^1, \dots, P^k occur at most r times, the corresponding formula $\chi_J(\bar{z}, Z)$ in 8.2.3 is existential, too, and contains Z at most r times, and $(\Pi, P^i) \bar{t}$ is equivalent to

$$(*) \quad \exists u[\text{LFP}_{\bar{z}, Z} \chi_J] \bar{t} \tilde{u}^1$$

By 8.7.3 this formula is equivalent to a formula in $\text{FO}(\text{BFP}_r)$.

Now in any S-DATALOG_r program of breadth $\leq r$ of the form $\Sigma := \Pi, \Pi_1$ we can deal with Π_1 as with Π , where for the intentional symbols of Π occurring in Π_1 (as extensional symbols!) we use the defining formulas in $\text{FO}(\text{BFP}_r)$ just gotten. For an arbitrary S-DATALOG program $\Sigma := \Pi_0, \dots, \Pi_n$ of breadth $\leq r$ the result follows by induction on n . \square

Using 8.7.1 we get:

Corollary 9.1.3 $\text{S-DATALOG}_0 \equiv \text{FO}$, $\text{S-DATALOG}_1 \equiv \text{FO}(\text{TC})$, and $\text{S-DATALOG}_2 \equiv \text{S-DATALOG}$. \square

What is the logic corresponding to plain DATALOG? We get a hint from the observation that in the formula in (*) of the preceding proof the least fixed-point operator is applied to an existential formula. In fact, we have:

Theorem 9.1.4 DATALOG and Existential Fixed-Point Logic $\text{E}(\text{LFP})$ have the same expressive power.

Here the class of formulas of $\text{E}(\text{LFP})$ is given by

- $\frac{}{\varphi}$ where φ is an atomic second-order formula
- $\frac{}{\neg\varphi}$ where φ is an atomic first-order formula
- $\frac{\varphi, \psi}{(\varphi \wedge \psi)}$, $\frac{\varphi, \psi}{(\varphi \vee \psi)}$, $\frac{\varphi}{\exists x\varphi}$
- $\frac{\varphi}{[\text{LFP}_{\bar{y}, Y} \varphi] \bar{t}}$

Proof (of 9.1.4). $\text{DATALOG} \leq \text{E}(\text{LFP})$: We saw in the proof of the preceding theorem that any DATALOG formula $(\Pi, P^i) \bar{t}$ is equivalent to the formula in (*). But this formula is a formula of $\text{E}(\text{LFP})$.

$\text{E}(\text{LFP}) \leq \text{DATALOG}$: The proof parallels that of $\text{FO}(\text{BFP}) \leq \text{S-DATALOG}$, but now in every step of the induction the S-DATALOG program given there has to be joined to a single DATALOG program. Proceed as follows:

¹ This formula is only equivalent to the DATALOG program in models with at least two elements. To remove this restriction one can proceed as follows: On structures of cardinality 1, the k -th stage is already the final one. But the k -th stage of P^1 can be expressed by an existential first-order formula ψ (compare 9.3.1). Then, for arbitrary structures, we have to replace the formula χ_J in (*) by $(\psi \vee \chi_J)$.

- For $\varphi(\bar{y})$ atomic set $\Pi_\varphi := \{Q_\varphi\bar{y} \leftarrow \varphi\}$.
- For $\varphi(\bar{y}) = \neg\psi$ with atomic ψ set $\Pi_\varphi := \{Q_\varphi\bar{y} \leftarrow \neg\psi\}$.
- For $\varphi(\bar{y}) = (\psi \wedge \chi)$ set $\Pi_\varphi := \Pi_\psi \cup \Pi_\chi \cup \{Q_\varphi\bar{y} \leftarrow Q_\psi\bar{y}, Q_\chi\bar{y}\}$.
- For $\varphi(\bar{y}) = (\psi \vee \chi)$ set $\Pi_\varphi := \Pi_\psi \cup \Pi_\chi \cup \{Q_\varphi\bar{y} \leftarrow Q_\psi\bar{y}, Q_\varphi\bar{y} \leftarrow Q_\chi\bar{y}\}$.
- For $\varphi(\bar{y}) = \exists x\psi(\bar{y}, x)$ set $\Pi_\varphi := \Pi_\psi \cup \{Q_\varphi\bar{y} \leftarrow Q_\psi\bar{y}x\}$.
- Suppose $\varphi(\bar{y}) = [\text{LFP}_{\bar{x}, X}\psi]\bar{t}$. Assume w.l.o.g. that the individual variables free in ψ are among \bar{x} . Consider X in ψ as a relation symbol (and not as a relation variable) and let Π_ψ be the DATALOG program corresponding to ψ by induction hypothesis and having X as an extensional symbol; X does not occur negated in the body of a clause, since ψ is positive in X . Then as Π_φ we can take the following program (that contains X as an intentional symbol): $\Pi_\varphi := \Pi_\psi \cup \{X\bar{x} \leftarrow Q_\psi\bar{x}, Q_\varphi\bar{y} \leftarrow X\bar{t}\}$. \square

Since in graphs the E(LFP) formula

$$[\text{LFP}_{xy, X} (Exy \vee \exists z(Xxz \wedge Ezy))] cd$$

expresses that there is a path from c to d , we see that $\text{E(LFP)} \not\leq \text{FO}$. On the other hand, we also have $\text{FO} \not\leq \text{E(LFP)}$. An example of an FO-formula not equivalent to a formula of E(LFP) is given by $\forall xPx$, since a simple proof by induction shows that E(LFP) has the following monotonicity property:

Proposition 9.1.5 *If $\varphi(x_1, \dots, x_n)$ is a formula of E(LFP) and $a_1, \dots, a_n \in \mathcal{A}$, then*

$$\mathcal{A} \models \varphi[a_1, \dots, a_n] \text{ and } \mathcal{A} \subseteq \mathcal{B} \text{ imply } \mathcal{B} \models \varphi[a_1, \dots, a_n]. \quad \square$$

Let $\tau \supseteq \{<, S, \min, \max\}$. In presence of S, \min, \max we have for ordered structures:

$$\mathcal{A} \subseteq \mathcal{B} \text{ implies } \mathcal{A} = \mathcal{B}.$$

Thus, on ordered structures, the monotonicity property of the preceding proposition holds for all logics.

Theorem 9.1.6 *Let $\tau \supseteq \{<, S, \min, \max\}$. On ordered structures, we have $\text{E(LFP)} \equiv \text{FO(LFP)}$, that is, the expressive power of DATALOG, S-DATALOG, E(LFP), FO(BFP), and FO(LFP) coincide.*

Proof. To show that $\text{FO(LFP)} \leq \text{E(LFP)}$ recall (cf. 8.2.9) that every formula of FO(LFP) is equivalent to a formula of the form

$$\exists x [\text{LFP}_{\bar{y}, Y}\psi]\bar{t}$$

where ψ is first-order. Thus it suffices to show that on ordered structures universal quantification can be replaced by an LFP operation (according to E(LFP)). For this note that $\forall x\varphi(x)$ is equivalent to

$$[\text{LFP}_{x, X} (\varphi(\min) \wedge x = \min) \vee \exists y(Xy \wedge Syx \wedge \varphi(x))] \max. \quad \square$$

Exercise 9.1.7 A DATALOG formula $(\Pi, P)\bar{t}$ is *positive* if no negated atomic formula at all occurs in the body of any clause. Introduce positive existential fixed-point logic and show that it has the same expressive power as positive DATALOG. \square

Exercise 9.1.8 Let $\Sigma := \Pi_0, \dots, \Pi_k$ and $\Sigma' := \Pi'_0, \dots, \Pi'_l$ be S-DATALOG programs with $\Pi_0 \cup \dots \cup \Pi_k = \Pi'_0 \cup \dots \cup \Pi'_l$. Show that $\mathcal{A}[\Sigma] = \mathcal{A}[\Sigma']$ for any structure \mathcal{A} . Hint: Consider a finest subdivision of $\Pi_0 \cup \dots \cup \Pi_k$ leading to an S-DATALOG program and show that Σ and Σ' are equivalent to it. \square

Exercise 9.1.9 Show that on ordered structures the successor relation S is not definable in E(LFP) in terms of $<, \min, \max$. (Hint: Use 9.1.5.) Show that $<$ is definable in E(LFP) using S . \square

Exercise 9.1.10 Extend the semantics of E(LFP) and DATALOG to infinite structures and show the equivalence of E(LFP) and DATALOG for arbitrary structures. Conclude

- If $\varphi(\bar{x})$ is a formula of E(LFP), \mathcal{A} is an arbitrary structure, $\bar{a} \in A$, and $\mathcal{A} \models \varphi[\bar{a}]$, then there is a finite $A_0 \subseteq A$ such that $\bar{a} \in A_0$, and for all \mathcal{B} with $A_0 \subseteq B$, if A_0 generates the same substructure in \mathcal{B} as in \mathcal{A} then $\mathcal{B} \models \varphi[\bar{a}]$.
- If an E(LFP)-formula has a model then it has a finite model. \square

9.2 I-DATALOG and P-DATALOG

Until now we have introduced some natural classes of general logic programs, and we have seen that they correspond to certain fixed-point logics. We now ask the other way round: What classes of programs correspond to, say, inflationary (= least) fixed-point logic and to partial fixed-point logic?

First we show that inflationary fixed-point logic FO(IFP) can be captured by general logic programs, if we equip them with an “inflationary” semantics. We speak of I-DATALOG, the “I” standing for the inflationary aspect of its semantics.

Being a general logic program, an I-DATALOG program Π may have clauses whose body also contains negated intentional symbols. The semantics is fixed in a way that consistently extends the conventions for DATALOG. Similarly as with DATALOG, we may suppose that for any intentional symbol P of Π there are distinct variables \bar{x}_P such that any clause in Π with head symbol P has the form

$$P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l.$$

We set

$$\varphi_P := \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \}$$

where \bar{y} contains the variables of the corresponding clause that are distinct from the variables in \bar{x}_P . Assume that P^1, \dots, P^k are the intentional symbols of Π . Then the simultaneous inflationary operations corresponding to $\varphi_{P^1}, \dots, \varphi_{P^k}$ give rise to the sequences

$$P_{(0)}^i := \emptyset, \quad P_{(n+1)}^i := P_{(n)}^i \cup \{\bar{x}_{P^i} \mid \varphi_{P^i}(\bar{x}_{P^i}, P_{(n)}^1, \dots, P_{(n)}^k)\}$$

with the simultaneous fixed-point given by $P_{(\infty)}^i := \bigcup_{n \geq 0} P_{(n)}^i$.

Let \mathcal{A} be a τ_{ext} -structure. Then Π leads to the τ -structure

$$\mathcal{A}[\Pi] := (\mathcal{A}, P_{(\infty)}^1, \dots, P_{(\infty)}^k).$$

If P is an intentional symbol of Π , $(\Pi, P)\bar{t}$ is an I-DATALOG *formula*, its meaning in a τ_{ext} -structure \mathcal{A} being $\bar{t} \in P^{\mathcal{A}[\Pi]}$.

Exercise 9.2.1 Let $\Pi = \{Px \leftarrow Rx, Qx \leftarrow \neg Px\}$ be an I-DATALOG program and Σ be the S-DATALOG program $\{Px \leftarrow Rx\}, \{Qx \leftarrow \neg Px\}$. Show for $\mathcal{A} = (A, R^A)$ that $Q^{\mathcal{A}[\Pi]} = A$ and $Q^{\mathcal{A}[\Sigma]} = A \setminus R^A$. In this sense the semantics of I-DATALOG and S-DATALOG programs are not compatible. \square

Theorem 9.2.2 I-DATALOG \equiv FO(IFP).

Proof. By the definitions just given, an I-DATALOG formula $(\Pi, P)\bar{t}$ with intentional symbols $(P =)P^1, P^2, \dots, P^k$ is equivalent to

$$[\text{S-IFP}_{\bar{x}_{P^1}, P^1, \dots, \bar{x}_{P^k}, P^k} \varphi_{P^1}, \dots, \varphi_{P^k}] \bar{t}$$

and hence, is definable in FO(IFP). Thus, I-DATALOG \leq FO(IFP).

To show FO(IFP) \leq I-DATALOG, to every formula φ of FO(IFP) with free first-order variables among \bar{y} we assign an equivalent I-DATALOG-formula $(\Pi_\varphi, Q_\varphi)\bar{y}$. Free relation variables in φ are treated as relation symbols, hence, they are extensional symbols of Π_φ . Π_φ contains an intentional 0-ary symbol DONE_φ . DONE_φ becomes true at the last step of the evaluation of Π_φ (more precisely, when DONE_φ is true, all intentional symbols have reached their final value). We proceed by induction on φ :

$$\varphi(\bar{y}) \text{ atomic: Set } \Pi_\varphi := \left\{ \begin{array}{l} Q_\varphi \bar{y} \leftarrow \varphi \\ \text{DONE}_\varphi \leftarrow \end{array} \right\}$$

$\varphi(\bar{y}) := \neg\psi$: Let Π_φ consist of the clauses of Π_ψ and

$$\begin{array}{l} Q_\varphi \bar{y} \leftarrow \text{DONE}_\psi, \neg Q_\psi \bar{y} \\ \text{DONE}_\varphi \leftarrow \text{DONE}_\psi \end{array}$$

$\varphi(\bar{y}) := (\psi \vee \chi)$: Let Π_φ consist of the clauses of $\Pi_\psi \cup \Pi_\chi$ and the clauses

$$\begin{array}{lcl}
 Q_{\varphi}\bar{y} & \leftarrow & Q_{\psi}\bar{y} \\
 Q_{\varphi}\bar{y} & \leftarrow & Q_{\chi}\bar{y} \\
 \text{DONE}_{\varphi} & \leftarrow & \text{DONE}_{\psi}, \text{DONE}_{\chi}
 \end{array}$$

$\varphi(\bar{y}) := \exists x\psi(\bar{y}, x)$: Let Π_{φ} consist of the clauses of Π_{ψ} and the clauses

$$\begin{array}{lcl}
 Q_{\varphi}\bar{y} & \leftarrow & Q_{\psi}\bar{y}x \\
 \text{DONE}_{\varphi} & \leftarrow & \text{DONE}_{\psi}
 \end{array}$$

$\varphi(\bar{y}) := [\text{IFP}_{\bar{x}, X}\psi]\bar{t}$: Assume w.l.o.g. that all first-order variables free in ψ are among $\bar{x} = x_1 \dots x_k$. By induction hypothesis there is a program Π_{ψ} corresponding to ψ (with X as an extensional symbol). Thus, for a given interpretation of X , the program Π_{ψ} returns as Q_{ψ} the set $\{\bar{x} \mid \psi(\bar{x}, X)\}$. For $\varphi(\bar{y})$ the program Π_{ψ} must be repeated several times until the fixed-point of the inflationary operation is reached. But note that in order to get the stages of the IFP operation, after each iteration step the intentional symbols of Π_{ψ} have to get empty again before the next iteration step starts. However, such a set back is excluded by the ‘‘inflationary’’ semantics of I-DATALOG. To overcome this difficulty we increase the arity of the intentional symbols, using the additional components for timestamps, the timestamps for the $(n + 1)$ -th iteration step being given by the new elements of X obtained in the n -th iteration step.

The program Π_{φ} given below contains as intentional symbols the intentional symbols of Π_{ψ} and X , Q_{φ} , DONE_{φ} , the 0-ary symbols NOTEMPTY , DELAY , DONE , the k -ary symbols OLD , I , and for every intentional symbol R of Π_{ψ} a new symbol R^* with $\text{arity}(R^*) = \text{arity}(R) + k$. We use the following notation: For new variables $\bar{z} = z_1 \dots z_k$, the program $\Pi_{\psi}^*[\bar{z}]$ is obtained by replacing in every clause of Π_{ψ} each (possibly negated) atomic formula of the form $R\bar{s}$ with an intentional R by $R^*\bar{s}\bar{z}$. And if Π is any program then $\Pi \parallel \gamma_1, \dots, \gamma_m$ is obtained by appending $\gamma_1, \dots, \gamma_m$ to the bodies of all clauses in Π .

The following remarks should help to read Π_{φ} : The clauses in (1) take care of the first iteration of ψ , (2) - (4) stop the program in case the first iteration returns an empty result, and (5) updates the value of X . The subsequent iterations of ψ are simulated by the clauses in (6) timestamped with the tuples added at the previous stage. By (7) the value of X is updated and the old value is stored in OLD by use of (8). (9) checks whether no new value has been added to X . In this case (10) and (12) stop the program. By (11) the goal symbol Q_{φ} gets its right value.

- (1) $\Pi_{\psi} \parallel \neg\text{DONE}_{\psi}$
- (2) $\text{NOTEMPTY} \leftarrow \text{DONE}_{\psi}, Q_{\psi}\bar{x}$
- (3) $\text{DELAY} \leftarrow \text{DONE}_{\psi}$
- (4) $\text{DONE}_{\varphi} \leftarrow \text{DELAY}, \neg\text{NOTEMPTY}$
- (5) $X\bar{x} \leftarrow \text{DELAY}, Q_{\psi}\bar{x}$

- (6) $\Pi_{\psi}^*[\bar{z}] \parallel X\bar{z}, \neg \text{OLD}\bar{z}$
- (7) $X\bar{x} \leftarrow \text{DONE}_{\psi}^*\bar{z}, Q_{\psi}^*\bar{x}\bar{z},$
- (8) $\text{OLD}\bar{z} \leftarrow \text{DONE}_{\psi}^*\bar{z}$
- (9) $I\bar{z} \leftarrow \text{DONE}_{\psi}^*\bar{z}, Q_{\psi}^*\bar{x}\bar{z}, \neg X\bar{x}$
- (10) $\text{DONE} \leftarrow \text{OLD}\bar{z}, \neg I\bar{z}$
- (11) $Q_{\varphi}\bar{y} \leftarrow \text{DONE}, X\bar{t}$
- (12) $\text{DONE}_{\varphi} \leftarrow \text{DONE}$ □

Remark 9.2.3 Recall the definition of FO(IFP, #), fixed-point logic with counting, given in subsection 8.4.1. DATALOG(#), *Datalog with Counting*, extends DATALOG by allowing two sorts (the point sort and the number sort) and admitting two-sorted intentional relation symbols and certain terms of the number sort in its clauses. More precisely: Call a term t of the number sort *basic* if it is either a variable μ or \min, \max or of the form $\#_x R\bar{y}\bar{\mu}$ for an intentional R . Now, a DATALOG(#) program is a finite set of clauses

$$\gamma \leftarrow \gamma_1, \dots, \gamma_l$$

where γ is atomic and relational and $\gamma_1, \dots, \gamma_l$ are atomic or negated atomic formulas, all of whose terms of the second sort are basic; moreover, intentional symbols do not occur in negated atomic formulas. This last condition is not a real restriction, since we already remarked in part (e) of 8.4.17 that $\neg R\bar{y}\bar{\mu}$ is equivalent to $\#_x R\bar{y}\bar{\mu} = \min$. So negated intentional relation symbols are available via unnegated atomic formulas containing counting terms. We equip DATALOG(#) with the inflationary semantics. Now it is not difficult to adapt the proof of the preceding theorem in order to show

$$\text{DATALOG}(\#) \equiv \text{FO}(\text{IFP}, \#). \quad \square$$

We turn to partial fixed-point logic and introduce P-DATALOG that will correspond to FO(PFP). The syntax of P-DATALOG is the same as that of I-DATALOG, that is, every general logic program is a P-DATALOG program. The semantics is given by the following conventions: Let Π be a P-DATALOG program. As in the preceding cases we assume that every clause with the intentional symbol P in its head has the form

$$P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l.$$

Again, set

$$\varphi_P(\bar{x}_P) := \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \}.$$

Assume that P^1, \dots, P^k are the intentional symbols of Π . Set

$$P_{(0)}^i := \emptyset, \quad P_{(n+1)}^i := \{ \bar{x}_{P^i} \mid \varphi_{P^i}(\bar{x}_{P^i}, P_{(n)}^1, \dots, P_{(n)}^k) \}$$

and let $(P_{(\infty)}^1, \dots, P_{(\infty)}^k)$ be the simultaneous fixed-point. Thus, in case there is an n_0 such that $P_{(n_0+1)}^i = P_{(n_0)}^i$ for $i = 1, \dots, k$, we have $P_{(\infty)}^1 = P_{(n_0)}^1, \dots, P_{(\infty)}^k = P_{(n_0)}^k$; otherwise, $P_{(\infty)}^i = \emptyset$ for $i = 1, \dots, k$. Let \mathcal{A} be a τ_{ext} -structure. Then Π leads to the τ -structure

$$\mathcal{A}[\Pi] := (\mathcal{A}, P_{(\infty)}^1, \dots, P_{(\infty)}^k).$$

For any intentional symbol Q of Π and terms \bar{t} , $(\Pi, Q)\bar{t}$ is a P-DATALOG formula, its meaning in a τ_{ext} -structure \mathcal{A} being $\bar{t} \in Q^{\mathcal{A}[\Pi]}$.²

Theorem 9.2.4 P-DATALOG \equiv FO(PFP).

Proof. By the definitions just given, a P-DATALOG formula $(\Pi, Q)\bar{t}$, where Π has the intentional symbols $(Q =)P^1, P^2, \dots, P^k$, is equivalent to

$$[\text{S-PFP}_{\bar{x}_{P^1}, P^1, \dots, \bar{x}_{P^k}, P^k} \varphi_{P^1}, \dots, \varphi_{P^k}] \bar{t}$$

and hence, to a formula of FO(PFP).

For the direction $\text{FO(PFP)} \leq \text{P-DATALOG}$ we proceed as in the case of $\text{FO(IFP)} \leq \text{I-DATALOG}$, assigning to every formula φ of FO(PFP) with free first-order variables among \bar{y} an equivalent P-DATALOG formula $(\Pi_\varphi, Q_\varphi)\bar{y}$, where Π_φ contains a 0-ary symbol DONE_φ taking the value TRUE in the last stage of the evaluation. Again, free second-order variables of φ are treated as extensional symbols in Π_φ . The atomic case and the cases $\varphi = \neg\psi$, $\varphi = (\psi \vee \chi)$ and $\varphi = \exists x\psi$ are handled as in 9.2.2. For $\varphi(\bar{y}) := [\text{PFP}_{\bar{x}, X\psi}] \bar{t}$, where w.o.l.g. the first-order variables free in ψ are among \bar{x} , the program Π_φ given below is simpler than the corresponding one for I-DATALOG, since in the semantics of P-DATALOG an intentional symbol Z gets \emptyset in an evaluation step in case no positive Z -information is obtained. Π_φ contains as intentional symbols the intentional symbols of Π_ψ , X , DONE_φ , Q_φ , and the 0-ary symbols UPDATE, NOTFIXPOINT.

The clauses in (1) take care of an application of Π_ψ . Clauses (3) and (4) check whether a fixed-point has been reached. In the positive case, clauses (7)–(10) stop the program. Clauses (5) and (6) serve to update the values of the intentional symbols for the next application of Π_ψ .

- (1)
$$\begin{array}{l} \Pi_\psi \\ X\bar{x} \leftarrow X\bar{x} \end{array} \parallel \neg\text{DONE}_\psi, \neg\text{UPDATE}$$
- (2) $\text{UPDATE} \leftarrow \text{DONE}_\psi$
- (3) $\text{NOTFIXPOINT} \leftarrow \text{DONE}_\psi, Q_\psi\bar{x}, \neg X\bar{x}$
- (4) $\text{NOTFIXPOINT} \leftarrow \text{DONE}_\psi, X\bar{x}, \neg Q_\psi\bar{x}$
- (5) $Q_\psi\bar{x} \leftarrow \text{DONE}_\psi, Q_\psi\bar{x}$
- (6) $X\bar{x} \leftarrow \text{UPDATE}, \text{NOTFIXPOINT}, Q_\psi\bar{x}$

² One easily verifies that the semantics extends the conventions for DATALOG.

- (7) $\text{DONE}_\varphi \leftarrow \text{UPDATE}, \neg\text{NOTFIXPOINT}$
- (8) $Q_\varphi \bar{y} \leftarrow \text{UPDATE}, \neg\text{NOTFIXPOINT}, Q_\psi \bar{t}$
- (9) $\text{UPDATE} \leftarrow \text{UPDATE}, \neg\text{NOTFIXPOINT}$
- (10) $Q_\psi \bar{x} \leftarrow \text{UPDATE}, \neg\text{NOTFIXPOINT}, Q_\psi \bar{x}$ □

9.3 A Preservation Theorem

A DATALOG program is *positive* if no negated atomic formula at all occurs in the body of any clause. Recall that, given a DATALOG program Π , an intentional symbol P , and $n \geq 0$, we denote by $P_{(n)}$ the n -th stage of P in the evaluation of Π .

Proposition 9.3.1 *Let Π be a DATALOG program. Then, for any intentional symbol P and $n \geq 0$, the relation $P_{(n)}$ is expressible by an existential first-order formula φ_P^n . If Π is positive then φ_P^n can be chosen existential positive (cf. 2.5.3).*

Proof. The proof proceeds by induction on n . If $n = 0$ then $\varphi_P^0 := F$. For $n > 0$, let

$$\varphi_P(\bar{x}_P) = \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \},$$

where \bar{y} contains the variables in $\gamma_1, \dots, \gamma_l$ that are not in \bar{x}_P . To obtain $\varphi_P^n(\bar{x}_P)$, replace in $\varphi_P(\bar{x}_P)$ any subformula $Q\bar{t}$ with intentional Q by $\varphi_Q^{n-1}(\bar{t})$. □

A DATALOG formula $(\Pi, Q)\bar{t}$ is *bounded* if there is an $n \geq 0$ such that $Q_{(n)} = Q_{(\infty)}$ holds in all structures. Thus a bounded (positive) DATALOG formula is equivalent to an existential (positive) first-order formula. On the other hand, an existential (positive) first-order formula is equivalent to a bounded (positive) DATALOG formula. In fact, let $\varphi(\bar{x})$ be existential (positive), w.l.o.g. of the form

$$\bigvee_{i=1, \dots, s} \exists y_1 \dots \exists y_{m_i} (\varphi_{i1} \wedge \dots \wedge \varphi_{ik_i}),$$

where the φ_{ij} are atomic or negated atomic (atomic) formulas. Then $\varphi(\bar{x})$ is equivalent to the DATALOG formula $(\Pi, Q)\bar{x}$, where Π has the clauses

$$Q\bar{x} \leftarrow \varphi_{i1}, \dots, \varphi_{ik_i}$$

for $i = 1, \dots, s$ and where Q is a symbol not in φ . As $Q_{(1)} = Q_{(\infty)}$, $(\Pi, Q)\bar{x}$ is bounded.

Thus, by the proposition, the *bounded* positive DATALOG formulas correspond to the existential positive first-order formulas. The following preservation theorem shows that they are the only positive DATALOG formulas equivalent to first-order formulas, thus documenting once more that first-order logic lacks the ability to express recursion.

Theorem 9.3.2 *For a positive DATALOG formula $(\Pi, Q)\bar{t}$ the following are equivalent:*

- (i) $(\Pi, Q)\bar{t}$ is equivalent to a first-order formula.
- (ii) $(\Pi, Q)\bar{t}$ is bounded.

At the end of this section we sketch a proof of this result. First we draw some consequences.

Corollary 9.3.3 *Suppose that the positive DATALOG formulas $(\Pi_1, Q_1)\bar{s}$ and $(\Pi_2, Q_2)\bar{t}$ are equivalent. If $(\Pi_1, Q_1)\bar{s}$ is bounded, then $(\Pi_2, Q_2)\bar{t}$ is bounded, too. \square*

Proposition 9.3.4 *For a positive DATALOG formula $(\Pi, Q)\bar{x}$ the following are equivalent:*

- (i) $(\Pi, Q)\bar{x}$ is bounded.
- (ii) $(\Pi, Q)\bar{x}$ is equivalent to an existential positive first-order formula.
- (iii) $(\Pi, Q)\bar{x}$ is equivalent to a first-order formula.
- (iv) *There is an $n \geq 1$ such that for all \mathcal{A} and $\bar{a} \in Q^{A^{[\Pi]}}$ there is a substructure \mathcal{B} of \mathcal{A} with $\|B\| \leq n$ and $\bar{a} \in Q^{B^{[\Pi]}}$.*

Proof. The equivalence of (i) and (ii) has already been shown, and that of (i) and (iii) is given by the theorem. Next we show that (ii) implies (iv): Assume (ii) and let $(\Pi, Q)\bar{x}$ be equivalent to the existential formula $\varphi(\bar{x})$, say $\varphi(\bar{x}) = \exists \bar{y}\psi(\bar{x}, \bar{y})$ with quantifier-free ψ . Given \mathcal{A} and $\bar{a} \in Q^{A^{[\Pi]}}$, that is, $\mathcal{A} \models \varphi[\bar{a}]$, choose \bar{b} such that $\mathcal{A} \models \psi[\bar{a}, \bar{b}]$. Then, for the substructure \mathcal{B} generated by $\{\bar{a}, \bar{b}\}$, we have $\mathcal{B} \models \varphi[\bar{a}]$ and hence, $\bar{a} \in Q^{B^{[\Pi]}}$. Thus, as n we can take the sum of $\text{length}(\bar{x})$, $\text{length}(\bar{y})$, and of the number of constant symbols.

Finally we show that (iv) implies (i). Clearly, given any n and any DATALOG formula $(\Pi, Q)\bar{x}$, there is an s such that $Q_{(s)} = Q_{(\infty)}$ in all structures of cardinality $\leq n$. Suppose that the positive DATALOG formula $(\Pi, Q)\bar{x}$ satisfies (iv). Choose n according to (iv) and choose s corresponding to $(\Pi, Q)\bar{x}$ and n . We show that $(\Pi, Q)\bar{x}$ is bounded by s . If \mathcal{A} and $\bar{a} \in Q^{A^{[\Pi]}}$ are given, (iv) yields a substructure \mathcal{B} of \mathcal{A} with $\|B\| \leq n$ such that $\bar{a} \in Q^{B^{[\Pi]}}$; therefore, $\bar{a} \in Q_{(s)}^{B^{[\Pi]}}$ and hence, $\bar{a} \in Q_{(s)}^{A^{[\Pi]}}$. \square

We give a further application of 9.3.2. A *universal Horn sentence* is a conjunction of sentences of the form

- (1) $\forall \bar{x}\psi$
- (2) $\forall \bar{x}((\psi_0 \wedge \dots \wedge \psi_m) \rightarrow \psi)$
- (3) $\forall \bar{x}(\neg\psi_0 \vee \dots \vee \neg\psi_m)$,

where the ψ 's and the ψ_i 's are atomic. A class K of τ -structures is *projective Horn* if there is a relational vocabulary σ disjoint from τ and a universal Horn sentence φ of vocabulary $\tau \cup \sigma$ such that the following holds:

- in all subformulas of φ of type (1) or (2) the formula ψ is relational, its relation symbol being from σ
- K is the class of τ -reducts of models of φ (that is, K is the class of models of the second-order sentence $\exists P^1 \dots \exists P^s \varphi$ where $\sigma = \{P^1, \dots, P^s\}$).

Since the first-order formulas involved are universal, a projective Horn class is closed under substructures. We have the following “compactness type” converse:

Theorem 9.3.5 *Let K be a projective Horn class. If K is axiomatizable in first-order logic then there is an $n \geq 1$ such that an arbitrary structure belongs to K if all substructures of cardinality at most n belong to K .*

Proof. Suppose σ, τ , and φ are as above and K is the class of τ -reducts of the models of φ . Choose a new 0-ary relation symbol Q and let Π be the positive existential DATALOG program obtained from φ by replacing each conjunct of the form (1) by the clause

$$(1^*) \quad \psi \leftarrow$$

each conjunct of the form (2) by the clause

$$(2^*) \quad \psi \leftarrow \psi_0, \dots, \psi_m$$

and each conjunct of the form (3) by the clause

$$(3^*) \quad Q \leftarrow \psi_0, \dots, \psi_m$$

together with

$$R\bar{x} \leftarrow R\bar{x}$$

for $R \in \sigma$. Hence, $\sigma \cup \{Q\}$ is the set of intentional symbols of Π . We show that the class of models of (Π, Q) is the complement K^c of K , that is, for every τ -structure \mathcal{A} ,

$$(+) \quad \mathcal{A} \in K^c \text{ iff } \mathcal{A} \models (\Pi, Q).$$

Then, K^c is axiomatizable by the positive DATALOG formula (Π, Q) and, as K , is first-order axiomatizable. Hence, K^c satisfies (iv) in 9.3.4, that is, there is an $n \geq 1$ such that all τ -structures $\mathcal{A} \in K^c$ contain a substructure of cardinality $\leq n$ belonging to K^c . Rephrasing this for K we get our claim.

To prove (+), assume first that $\mathcal{A} \in K^c$. Then, in particular, $\mathcal{B} := (\mathcal{A}, (P^{\mathcal{A}[\Pi]})_{P \in \sigma}) \not\models \varphi$. By its definition, \mathcal{B} satisfies the subformulas of type (1) and (2) in φ . Hence, there is a subformula of type (3), say $\forall \bar{x} (\neg \psi_0 \vee \dots \vee \neg \psi_m)$, such that $\mathcal{B} \models \neg \forall \bar{x} (\neg \psi_0 \vee \dots \vee \neg \psi_m)$. Choose $\bar{a} \in A$ such that $\mathcal{B} \models (\psi_0 \wedge \dots \wedge \psi_m)[\bar{a}]$. By (3*), $\mathcal{A} \models (\Pi, Q)$.

Conversely, suppose $\mathcal{A} \in K$, say $(\mathcal{A}, (P^{\mathcal{A}})_{P \in \sigma}) \models \varphi$. A simple induction on n shows that $P_{(\bar{n})} \subseteq P^{\mathcal{A}}$ for $P \in \sigma$; hence, $P^{\mathcal{A}[\Pi]} \subseteq P^{\mathcal{A}}$. Now, if we had $\mathcal{A} \models (\Pi, Q)$, then for some clause of the form (3*) and some $\bar{a} \in A$, we would have $(\mathcal{A}, (P^{\mathcal{A}[\Pi]})_{P \in \sigma}) \models \psi_0 \wedge \dots \wedge \psi_m[\bar{a}]$ and hence, $(\mathcal{A}, (P^{\mathcal{A}})_{P \in \sigma}) \models \psi_0 \wedge \dots \wedge \psi_m[\bar{a}]$ contrary to $(\mathcal{A}, (P^{\mathcal{A}})_{P \in \sigma}) \models \varphi$. \square

Theorem 9.3.5 can be used to prove nonaxiomatizability results. For example, the class K of acyclic digraphs is projective Horn as can be seen via the axiom

$$\exists P(\forall x\forall y(Exy \rightarrow Pxy) \wedge \forall x\forall y((Pxy \wedge Pyz) \rightarrow Pxz) \wedge \forall x\neg Pxx)$$

(take as P the transitive closure of E). Since K does not satisfy the conclusion in Theorem 9.3.5, acyclicity is not expressible in first-order logic.

We close by sketching the main idea of a proof of Theorem 9.3.2 using concepts and results from section 2.5 and refer to [7] for the full proof. Let (Π, Q) be a positive DATALOG formula where, for simplicity, we assume that Q is 0-ary. Suppose that (Π, Q) is equivalent to the first-order sentence φ . We have to show that (Π, Q) is bounded. Otherwise, (Π, Q) and, hence, φ have arbitrarily large minimal models. On the other hand, by positivity, (Π, Q) and, hence, φ are preserved under strict homomorphisms. By 2.5.5, for some l and m , no minimal model of φ has an l -scattered subset of cardinality $\geq m$. Finally, one obtains a contradiction by showing that in large minimal models there are such subsets.

9.4 Normal Forms for Fixed-Point Logics

The equivalence of I-DATALOG with FO(IFP) and of P-DATALOG with FO(PFP) (together with 8.2.3, the Lemma on Simultaneous Induction) provides an alternative way to derive normal forms for FO(IFP) and FO(PFP). We shall exemplify this first and then further improve these results, thus obtaining normal forms that, in particular, yield some kind of measure to compare FO(IFP) and FO(PFP) with transitive closure logic and bounded fixed-point logic. Moreover, they will enable us in Chapter 12 to get representations of FO(IFP) and FO(PFP) as logics with generalized quantifiers.

Once more, for simplicity, we consider only structures with at least two elements. In particular, we say that φ and ψ are equivalent if they are equivalent in all finite structures with at least two elements.³ And we say that φ and $\exists(\forall)x\psi$ are equivalent if φ is equivalent (in this sense) to both $\exists x\psi$ and $\forall x\psi$.

Theorem 9.4.1 *Let φ be an FO(IFP)-formula. Then there is an existential first-order formula ψ such that φ is equivalent to*

$$\exists(\forall)u[\text{IFP}_{\bar{z}, Z}\psi]\tilde{u}.$$

Proof. Let $\varphi = \varphi(\bar{v})$ with free first-order variables among \bar{v} . Choose an I-DATALOG formula $(\Pi, Q)\bar{v}$ equivalent to φ and let P^1, \dots, P^k with $Q = P^1$ be the intentional symbols of Π .

³ In most cases the restriction to structures with at least two elements can be removed as in similar situations before, or the corresponding statements are trivially false.

We may assume that the clauses of Π have the form

$$P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l$$

where \bar{x}_P is a sequence of distinct variables. Let

$$\varphi_P(\bar{x}_P) := \bigvee \{ \exists \bar{y} (\gamma_1 \wedge \dots \wedge \gamma_l) \mid P\bar{x}_P \leftarrow \gamma_1, \dots, \gamma_l \text{ in } \Pi \}$$

where \bar{y} contains all the variables of the corresponding clause not in \bar{x}_P . Then $(\Pi, Q)\bar{v}$ is equivalent to

$$[\text{S-IFP}_{\bar{x}_{P^1}, P^1, \dots, \bar{x}_{P^k}, P^k} \varphi_{P^1}, \dots, \varphi_{P^k}] \bar{v}.$$

Since $\varphi_{P^1}, \dots, \varphi_{P^k}$ are existential, the corresponding χ_J in 8.2.3 is existential, too, and

$$\models_{\text{fin}} [\text{S-IFP}_{\bar{x}_{P^1}, P^1, \dots, \bar{x}_{P^k}, P^k} \varphi_{P^1}, \dots, \varphi_{P^k}] \bar{v} \leftrightarrow \exists (\forall) u [\text{IFP}_{\bar{z}, Z} \chi_J] \bar{v} \bar{u}.$$

By 8.2.4, the formula on the right side is equivalent to a formula of the form $\exists (\forall) u [\text{IFP}_{\bar{z}, Z} \psi] \bar{u}$ with existential ψ . \square

The corresponding result for FO(PFP) can be obtained along the same lines, using $\text{FO(PFP)} \equiv \text{P-DATALOG}$. Moreover, we already proved it in 8.3.14 by different methods.

For FO(LFP) the kernel of the fixed-point operator in a normal form cannot be chosen existential (cf. the example before 9.1.5); however, we have

Theorem 9.4.2 *Suppose that φ is an FO(LFP)-formula. Then there is a Δ_2 -formula ψ , that is, a first-order formula ψ equivalent to both a Σ_2 -formula and a Π_2 -formula, such that φ is equivalent to*

$$\exists (\forall) u [\text{LFP}_{\bar{z}, Z} \psi] \bar{u}.$$

Proof. By the previous theorem there is an existential formula χ such that

$$\models_{\text{fin}} \varphi \leftrightarrow \exists (\forall) u [\text{IFP}_{\bar{z}, Z} \chi] \bar{u}.$$

By 8.2.11, for every existential χ there is a Δ_2 -formula ψ such that

$$\models_{\text{fin}} \exists (\forall) u [\text{IFP}_{\bar{z}, Z} \chi] \bar{u} \leftrightarrow \exists (\forall) u [\text{LFP}_{\bar{y}, Y} \psi] \bar{u}.$$

Thus, φ is equivalent to $\exists (\forall) u [\text{LFP}_{\bar{y}, Y} \psi] \bar{u}$. \square

Our next goal is to further simplify the normal forms we have obtained so far with respect to the structure of the kernel, in particular with respect to the form in which quantifiers appear in it. We first treat the case of FO(PFP) and FO(IFP). To state the result we need the notion of a nontrivial formula. An r -ary relation R on a set A is *nontrivial* if $\emptyset \neq R \neq A^r$. Furthermore, a formula of FO(PFP) is said to be *nontrivial* if for any subformula $\psi = [\text{PFP}_{\bar{x}, X} \chi] \bar{t}$ the stages F_1^x, F_2^x, \dots are nontrivial (for any assignment). The definitions for FO(IFP) and FO(LFP) are similar.

Theorem 9.4.3 *Let φ be an FO(PFP)-formula. Then φ is equivalent to a nontrivial, totally defined formula of the form*

$$\exists(\forall)u[\text{PFP}_{\bar{z},Z}(\psi_0 \vee \exists \bar{x} \in Z \exists \bar{y} \notin Z \psi_1)] \tilde{u}$$

where ψ_0 and ψ_1 are quantifier-free and do not contain Z . The same result is true if PFP is replaced by IFP everywhere.

We have seen in 8.7.1(c) that TC operations correspond to LFP operations in which every element of a new stage is already witnessed by a single element of the preceding stage. Similarly, the fixed-point operations of bounded fixed-point logic FO(BFP) correspond to LFP operations where every element of a new stage is witnessed by two elements (cf. 8.7.1(d)). The preceding theorem tells us that IFP operations and PFP operations can be rewritten in such a way that an element of a new stage is witnessed by two elements, but now one belonging to the preceding stage, the other one belonging to its complement.

The proof of 9.4.3 is given in three steps; each step essentially corresponds to one of the next lemmas. We state the lemmas and the proofs for FO(PFP). They remain literally true for FO(IFP) with a single exception which we point out at the corresponding place.

Lemma 9.4.4 *Let φ be an existential (or quantifier-free) first-order formula containing a second-order variable X . Then, provided X is interpreted by a nontrivial relation, φ is equivalent to a formula of the form*

$$\exists \bar{x}_1 \in X \dots \exists \bar{x}_k \in X \exists \bar{y}_1 \notin X \dots \exists \bar{y}_l \notin X \psi$$

where ψ is existential (or quantifier-free) and does not contain X . If φ is positive in X then l can be chosen to be 0.

Lemma 9.4.5 *Suppose that*

$$\varphi = \exists(\forall)u[\text{PFP}_{\bar{x},X}(\varphi_0 \vee \exists \bar{x}_1 \in X \dots \exists \bar{x}_k \in X \exists \bar{y}_1 \notin X \dots \exists \bar{y}_l \notin X \varphi_1)] \tilde{u}$$

is a nontrivial, totally defined formula, where φ_0 is quantifier-free, φ_1 is existential, and X does not occur in φ_0 and φ_1 . Then φ is equivalent to a nontrivial, totally defined formula of the form

$$\exists(\forall)u[\text{PFP}_{\bar{y},Y}(\psi_0 \vee \exists \bar{v} \in Y \exists \bar{w} \notin Y \psi_1)] \tilde{u}$$

where ψ_0 is quantifier-free, ψ_1 is existential, and Y does not occur in ψ_0, ψ_1 .

Lemma 9.4.6 *If φ is as in the conclusion of the preceding lemma then φ is equivalent to a formula of the same form with the additional property that ψ_1 is quantifier-free.*

Postponing the proof of these lemmas, we give the

Proof (of 9.4.3). Let φ be an FO(PFP)-formula. By 8.3.14 we can assume that

$$\varphi = \exists(\forall)u[\text{PFP}_{\bar{u},U}\chi]\bar{u},$$

that φ is totally defined, and that χ is existential. Suppose U is s -ary. First, we pass to a nontrivial formula. For this purpose we introduce an $(s + 2)$ -ary relation symbol Y where we use the additional components to ensure nontriviality by setting

$$\rho(\bar{u}, u_{s+1}, u_{s+2}) := \rho_0(\bar{u}, u_{s+1}, u_{s+2}) \vee \rho_1(\bar{u}, u_{s+1}, u_{s+2}, Y)$$

with

$$\begin{aligned} \rho_0(\bar{u}, u_{s+1}, u_{s+2}) &:= u_s = u_{s+1} \wedge \neg u_{s+1} = u_{s+2} \\ \rho_1(\bar{u}, u_{s+1}, u_{s+2}, Y) &:= u_{s+1} = u_{s+2} \wedge \exists \bar{v} v_{s+1} v_{s+2} Y \bar{v} v_{s+1} v_{s+2} \\ &\quad \wedge \chi(Y_{-} u_{s+1} u_{s+2}). \end{aligned}$$

The stages $F_1^\rho, F_2^\rho, \dots$ are nontrivial, since they contain all tuples of the form $\dots uvv$ with $u \neq v$ and no tuple of the form $\dots uvu$ with $u \neq v$. Moreover, $F_1^\rho = \{(u_1, \dots, u_{s+2}) \mid u_s = u_{s+1} \text{ and } u_{s+1} \neq u_{s+2}\}$. An evaluation of the stages shows

$$F_n^\chi = F_{n+1}^\rho - uu \text{ for } n \geq 0, \text{ and } F_\infty^\chi = F_\infty^\rho - uu.$$

By 9.4.4, provided a nontrivial relation is assigned to Y , ρ_1 is equivalent to a formula ψ_1 of the form

$$\psi_1 := \exists \bar{x}_1 \in Y \dots \exists \bar{x}_k \in Y \exists \bar{y}_1 \notin Y \dots \exists \bar{y}_l \notin Y \psi$$

where ψ is existential and does not contain Y . Using dummy variables if necessary, we may assume that $k \geq 1$. Then, if Y is interpreted by the empty relation, both ρ_1 and ψ_1 are false. Hence, we have $F_n^{\rho_0 \vee \rho_1} = F_n^{\rho_0 \vee \psi_1}$ for all n . Thus, φ is equivalent to the nontrivial, totally defined formula

$$\exists(\forall)u[\text{PFP}_{\bar{y},Y}(\rho_0(\bar{y}) \vee \exists \bar{x}_1 \in Y \dots \exists \bar{x}_k \in Y \exists \bar{y}_1 \notin Y \dots \exists \bar{y}_l \notin Y \psi)]\bar{u}.$$

As claimed, ρ_0 is quantifier-free, ψ is existential, and Y does not occur in ρ_0 and ψ . The theorem now follows by first applying Lemma 9.4.5 and then Lemma 9.4.6. \square

Proof (of 9.4.4). Let φ be an existential (or, quantifier-free) formula containing the variable X . We may assume that all negation symbols in φ are in front of atomic formulas. For such φ we inductively define φ^* of the form claimed. If $\varphi = X\bar{t}$ set $\varphi^* := \exists \bar{x} \in X \bar{t} = \bar{x}$. If $\varphi = \neg X\bar{t}$ set $\varphi^* := \exists \bar{y} \notin X \bar{t} = \bar{y}$. Let φ^* be φ for any other atomic or negated atomic formula φ . The definition for the remaining cases is straightforward, using the induction hypothesis and equivalences such as that of $(\varphi_1 \vee \exists \bar{x}_1 \in X \dots \exists \bar{x}_k \in X \exists \bar{y}_1 \notin X \dots \exists \bar{y}_l \notin X \varphi_2)$ and $\exists \bar{x}_1 \in X \dots \exists \bar{x}_k \in X \exists \bar{y}_1 \notin X \dots \exists \bar{y}_l \notin X (\varphi_1 \vee \varphi_2)$, provided a nontrivial relation is assigned to X and the variables in $\bar{x}_1, \dots, \bar{y}_l$ do not occur in φ_1 . \square

Proof (of 9.4.5). Suppose X is r -ary. The lemma claims that the number of quantifiers $\exists \bar{x} \in X, \exists \bar{y} \notin X$ in the kernel of the fixed-point operator can be reduced to a single positive and a single negative one. We achieve this by increasing the arity of the second-order variable to $(k + l) \cdot r$, whose first k r -blocks will consist of r -tuples in X and the last l r -blocks of r -tuples in the complement X^c of X . The corresponding r -tuples are incorporated in $(k + l)$ steps into the distinct components, thus making necessary additional components for a counter: One iteration step of X will correspond to $(k + l)$ steps for the new variable.

To simplify the presentation, assume that $k = l = 2$ and that X is unary, hence

$$\varphi = \exists(\forall)u[\text{PFP}_{x,X}(\varphi_0 \vee \exists x_1 \in X \exists x_2 \in X \exists y_1 \notin X \exists y_2 \notin X \varphi_1)] \tilde{u}$$

where $\varphi_1(x, x_1, x_2, y_1, y_2)$ is existential and $\varphi_0(x)$ and φ_1 do not contain X . Let X_n be the n -th stage of the operation corresponding to the PFP operator in φ . Choose a $(4 + 4)$ -ary new variable Y , where the first 4 components are used as a counter for the 4 steps corresponding to one step of X , the next two ones for elements of X , and the last two ones for elements of X^c . For the counters we introduce formulas $\epsilon_i(z_1, \dots, z_4)$ for $i = 1, \dots, 4$ by

$$\begin{aligned} \epsilon_1(z_1, \dots, z_4) &:= z_1 = z_2 = z_3 = z_4 \\ \epsilon_2(z_1, \dots, z_4) &:= \neg z_1 = z_2 \wedge z_1 = z_3 = z_4 \\ \epsilon_3(z_1, \dots, z_4) &:= \neg z_1 = z_3 \wedge z_1 = z_2 = z_4 \\ \epsilon_4(z_1, \dots, z_4) &:= \neg z_1 = z_4 \wedge z_1 = z_2 = z_3. \end{aligned}$$

If Y_n denotes the n -th stage of the operation corresponding to the PFP operator in the formula ψ below, set

$$Y_n i_- := \{ \bar{x} \mid \text{there are } \bar{z} \text{ such that } \epsilon_i(\bar{z}) \text{ and } Y_n \bar{z} \bar{x} \}$$

($\bar{x} = x_1 \dots x_4$ will correspond to $x_1 x_2 y_1 y_2$ in φ_1 above). For all $m \geq 0$ and any structure \mathcal{A} we will have (1) and (2).

$$(1) \quad \begin{aligned} Y_{4m+1} 1_- &= X_{m+1} \times A \times A \times A \\ Y_{4m+2} 2_- &= X_{m+1} \times A \times X_{m+1}^c \times A \\ Y_{4m+3} 3_- &= X_{m+1} \times X_{m+1} \times X_{m+1}^c \times A \\ Y_{4m+4} 4_- &= X_{m+1} \times X_{m+1} \times X_{m+1}^c \times X_{m+1}^c. \end{aligned}$$

Here the IFP case differs from the PFP case; the values for IFP are:

$$(1') \quad \begin{aligned} Y_{4m+1} 1_- &= \bigcup_{l \leq m+1} X_l \times A \times A \times A \\ Y_{4m+2} 2_- &= \bigcup_{l \leq m+1} X_l \times A \times X_l^c \times A \\ Y_{4m+3} 3_- &= \bigcup_{l \leq m+1} X_l \times X_l \times X_l^c \times A \\ Y_{4m+4} 4_- &= \bigcup_{l \leq m+1} X_l \times X_l \times X_l^c \times X_l^c. \end{aligned}$$

$$(2) \quad Y_{4m+i+j}i_- = Y_{4m+i}i_- \quad \text{for } 1 \leq i \leq 4 \text{ and } 1 \leq j \leq 3.$$

This is achieved if we use the formula $\psi :=$

$$\exists(\forall)u[\text{PFP}_{\bar{w}\bar{z},Y}(\epsilon_1(\bar{w}) \wedge \varphi_0(z_1)) \vee \exists\bar{u}\bar{x} \in Y \exists\bar{v}\bar{y} \notin Y \bigvee_{i=1}^4 (\epsilon_i(\bar{w}) \wedge \psi_i)] \bar{u}$$

where the ψ_i together with their meaning (formulated in a way suitable for an inductive proof of (1) and (2)) are given as follows:

$$\begin{aligned} \psi_1 &:= \epsilon_4(\bar{u}) \wedge \varphi_1(z_1, x_1, x_2, x_3, x_4) \\ &\quad \text{“}z_1 \in X_{m+1} \text{ is witnessed by } \bar{x} \in Y_{4m+4}_- \text{ via } \varphi_1(z_1, x_1, x_2, x_3, x_4)\text{”} \\ \psi_2 &:= \epsilon_1(\bar{u}) \wedge \epsilon_1(\bar{v}) \wedge x_2 = y_2 \wedge x_3 = y_3 \wedge x_4 = y_4 \wedge z_1 = x_1 \wedge z_3 = y_1 \\ &\quad \text{“}z_3 \in X_{m+1}^c \text{ is witnessed by } \bar{x} \in Y_{4m+1}1_- \text{ and } z_3x_2x_3x_4 \notin Y_{4m+1}1_- \text{”} \\ \psi_3 &:= \epsilon_2(\bar{u}) \wedge \epsilon_2(\bar{v}) \wedge x_1 = y_1 \wedge x_2 = y_2 \wedge x_4 = y_4 \wedge z_1 = x_1 \wedge z_2 = y_3 \\ &\quad \wedge z_3 = x_3 \\ &\quad \text{“}z_2 \in X_{m+1} \text{ is witnessed by } \bar{x} \in Y_{4m+2}2_- \text{ and } x_1x_2z_2x_4 \notin Y_{4m+2}2_- \text{”} \\ \psi_4 &:= \epsilon_3(\bar{u}) \wedge \epsilon_3(\bar{v}) \wedge x_2 = y_2 \wedge x_3 = y_3 \wedge x_4 = y_4 \wedge z_1 = x_1 \wedge z_2 = x_2 \\ &\quad \wedge z_3 = x_3 \wedge z_4 = y_1 \\ &\quad \text{“}z_4 \in X_{m+1}^c \text{ is witnessed by } \bar{x} \in Y_{4m+3}3_- \text{ and } z_4x_2x_3x_4 \notin Y_{4m+3}3_- \text{”}. \end{aligned}$$

Since X_∞ exists, Y_∞ exists, too, and $Y_\infty 1_- = X_\infty \times A \times A \times A$. As $\epsilon_1(u, u, u, u)$, but not $\epsilon_i(u, u, u, u)$ for $i = 2, 3, 4$, we thus see that φ and ψ are equivalent. \square

Proof (of 9.4.6). When proving this lemma, we again use additional components to simulate the unbounded existential quantifiers. Let

$$\varphi = \exists(\forall)u[\text{PFP}_{\bar{x},X}(\varphi_0 \vee \exists\bar{v} \in X \exists\bar{w} \notin X \varphi_1)] \bar{u}$$

with quantifier-free φ_0 and existential φ_1 , both not containing X . Assume that X is r -ary and

$$\varphi_1 = \exists y_1 \dots \exists y_l \chi(\bar{x}, \bar{v}, \bar{w}, \bar{y})$$

with quantifier-free χ . Let Z be $(r+l)$ -ary. Given a structure \mathcal{A} , let X_0, X_1, \dots denote the stages of the PFP operator in φ and Z_0, Z_1, \dots the stages of the PFP operator in

$$\psi = \exists(\forall)u[\text{PFP}_{\bar{x}\bar{y},Z}(\varphi_0 \vee \exists\bar{v}\bar{y}_1 \in Z \exists\bar{w}\bar{y}_2 \notin Z \chi(\bar{x}, \bar{v}, \bar{w}, \bar{y}_1))] \bar{u}.$$

Then a simple induction shows that $Z_n = X_n \times A^l$. Thus φ and ψ are equivalent. \square

Turning to FO(LFP), the improved normal form reads as follows:

Theorem 9.4.7 *Let φ be an FO(LFP)-formula. Then φ is equivalent to a nontrivial formula of the form*

$$(1) \quad \exists(\forall)u[\text{LFP}_{\bar{z},Z}(\psi_0 \vee \forall \bar{x}\exists \bar{y} \in Z \psi_1)] \tilde{u}$$

with quantifier-free ψ_0 and ψ_1 not containing Z , and φ is also equivalent to a formula of the form

$$(2) \quad \exists(\forall)u[\text{LFP}_{\bar{z},Z} \exists \bar{x}\forall \bar{y} \notin Z \psi_1] \tilde{u}$$

with quantifier-free ψ_1 not containing Z (note that $\forall \bar{y} \notin Z \psi_1$ abbreviates $\forall \bar{y}(Z\bar{y} \vee \psi_1)$ and hence, is positive in Z).

Moreover, in (1) we can choose ψ_1 of the form $\psi_{11} \rightarrow \psi_{12}$, where no variable of \bar{y} is free in ψ_{11} and no variable of \bar{z} is free in ψ_{12} .

Proof. We sketch a proof along the lines of the proof for the corresponding result 9.4.3 for FO(PFP). To obtain (1), we start from the fact (cf. 9.4.2) that φ is equivalent to a formula of the form

$$\exists(\forall)u[\text{LFP}_{\bar{u},U} \forall \bar{v}\exists \bar{w}\chi] \tilde{u}$$

where χ is quantifier-free (and positive in U). Let s be the arity of U . Using an $(s+2)$ -ary Y , we ensure nontriviality as in 9.4.3 by passing to

$$\rho(\bar{u}, u_{s+1}, u_{s+2}) := \rho_0(\bar{u}, u_{s+1}, u_{s+2}) \vee \rho_1(\bar{u}, u_{s+1}, u_{s+2}, Y)$$

with

$$\begin{aligned} \rho_0(\bar{u}, u_{s+1}, u_{s+2}) &:= u_s = u_{s+1} \wedge \neg u_{s+1} = u_{s+2} \\ \rho_1(\bar{u}, u_{s+1}, u_{s+2}, Y) &:= u_{s+1} = u_{s+2} \wedge \exists \bar{u}' u'_{s+1} u'_{s+2} Y \bar{u}' u'_{s+1} u'_{s+2} \\ &\quad \wedge \forall \bar{v}\exists \bar{w}\chi(Y_{-} u_{s+1} u_{s+2}). \end{aligned}$$

The formula ρ_1 is equivalent to

$$\forall \bar{v}\exists \bar{w}(u_{s+1} = u_{s+2} \wedge \exists \bar{u}' u'_{s+1} u'_{s+2} Y \bar{u}' u'_{s+1} u'_{s+2} \wedge \chi(Y_{-} u_{s+1} u_{s+2})).$$

As $\chi(Y_{-} u_{s+1} u_{s+2})$ is positive in Y , by 9.4.4 it is equivalent to

$$\exists \bar{y}_1 \in Y \dots \exists \bar{y}_k \in Y \chi_1$$

with quantifier-free χ_1 not containing Y , provided Y is interpreted by a non-trivial relation. Using dummy variables if necessary, we can assume that $k \geq 1$. Altogether, φ is equivalent to

$$\varphi_1 := \exists(\forall)u[\text{LFP}_{\bar{y},Y}(\rho_0 \vee \forall \bar{v}\exists \bar{w}\exists \bar{y}_1 \in Y \dots \exists \bar{y}_k \in Y \chi_2)] \tilde{u}$$

where $\rho_0(\bar{y})$ and $\chi_2(\bar{y}, \bar{v}, \bar{w}, \bar{y}_1, \dots, \bar{y}_k)$ are quantifier-free and do not contain Y . Let Z be $k \cdot r$ -ary where r is the arity of Y . Then for $n \geq 0$, if Y_n denotes the n -th stage of the LFP operation in φ_1 ,

$$Z_n := \underbrace{Y_n \times \dots \times Y_n}_{k \text{ times}}$$

are the stages of the LFP operations in

$$\exists(\forall)u[\text{LFP}_{\bar{z}_1 \dots \bar{z}_k, Z}((\rho_0(\bar{z}_1) \wedge \dots \wedge \rho_0(\bar{z}_k)))\forall$$

$$\forall \bar{u} \forall \bar{v} \exists \bar{w} \exists \bar{y}_1 \dots \bar{y}_k \in Z(\bigvee_{i=1}^k \bar{u} = \bar{z}_i \rightarrow (\rho_0(\bar{u}) \vee \chi_1(\bar{u}, \bar{v}, \bar{w}, \bar{y}_1, \dots, \bar{y}_k)))] \tilde{u}.$$

Now, as we have already done in previous proofs, we can eliminate the existential quantifiers $\exists \bar{w}$ by replacing Z , thereby using a relation variable of higher arity. We thus obtain a formula equivalent to φ of the form claimed in (1).

To obtain (2), we may assume (again by 9.4.2) that φ has the form

$$\varphi = \exists(\forall)u[\text{LFP}_{\bar{y}, Y} \exists \bar{v} \forall \bar{w} \psi] \tilde{u}$$

with quantifier-free ψ . Suppose that Y is r -ary. A simple induction, using the equivalence of $Y \bar{y}$ and $\forall \bar{z} \notin Y \neg \bar{z} = \bar{y}$, shows that ψ is equivalent to a formula

$$\forall \bar{z}_1 \notin Y \dots \forall \bar{z}_k \notin Y \rho(\bar{y}, \bar{v}, \bar{w}, \bar{z}_1, \dots, \bar{z}_k)$$

with quantifier-free ρ not containing Y , provided Y is interpreted by a relation different from the set of all r -tuples.

So we may assume that φ is equivalent to the formula

$$\varphi_1 := \exists(\forall)u[\text{LFP}_{\bar{y}, Y} \exists \bar{v} \forall \bar{w} \forall \bar{z}_1 \notin Y \dots \forall \bar{z}_k \notin Y \rho(\bar{y}, \bar{v}, \bar{w}, \bar{z}_1, \dots, \bar{z}_k)] \tilde{u}$$

(note that the proviso “ Y different from the set of all r -tuples” is irrelevant here, as the fixed-point operations in φ and φ_1 become stationary in case they arrive at the set of all r -tuples). Now, let Z be $k \cdot r$ -ary and let φ_2 be the formula

$$\exists(\forall)u[\text{LFP}_{\bar{y}_1 \dots \bar{y}_k, Z} \exists \bar{y} \exists \bar{v} \forall \bar{w} \forall \bar{z}_1 \dots \bar{z}_k \notin Z \bigvee_{i=1}^k (\bar{y} = \bar{y}_i \wedge \rho(\bar{y}, \bar{v}, \bar{w}, \bar{z}_1, \dots, \bar{z}_k))] \tilde{u}.$$

Then, for the complements Y_n^c and Z_n^c of the stages of the LFP operations in φ_1 and φ_2 , we have

$$Z_n^c = \underbrace{Y_n^c \times \dots \times Y_n^c}_{k \text{ times}}$$

Hence, φ_1 and φ_2 are equivalent. To eliminate the universal quantifiers $\forall \bar{w}$ in φ_2 , we use a variable X with $\text{arity}(X) = \text{length}(\bar{w}) + \text{arity}(Z)$ and pass to the equivalent formula

$$\begin{aligned} \exists(\forall)u[\text{LFP}_{\bar{y}, \bar{y}_1 \dots \bar{y}_k, X} \exists \bar{y} \exists \bar{v} \forall \bar{z}' \bar{z}_1 \dots \bar{z}_k \notin X \bigvee_{i=1}^k (\bar{y} = \bar{y}_i \\ \wedge \rho(\bar{y}, \bar{v}, \bar{z}', \bar{z}_1, \dots, \bar{z}_k))] \tilde{u}, \end{aligned}$$

a formula of the desired form (note that $X_n^c = \text{universe}^{\text{length}(\bar{w})} \times Z_n^c$ for all n). \square

As an application of the preceding theorem we treat a variant of the TC operator, the so-called *alternating transitive closure operator* ATC which turns out to be equivalent to the LFP operator.

We consider structures of the form $\mathcal{A} = (A, E^A, U^A)$ with binary E and unary U . In this context the elements of U^A are called *universal*, the elements of $A \setminus U^A$ *existential*. The *alternating transitive closure* $\text{ATC}(E^A, U^A)$ of (E^A, U^A) consists of those pairs $(a, b) \in A \times A$ for which there is an E^A -path from a to b with the property that whenever a universal point c is passed there must also be an E^A -path with this property from d to b for each $d \neq b$ with $E^A cd$. More precisely, using the LFP operator, we define:

$$(a, b) \in \text{ATC}(E^A, U^A) \text{ iff } (A, E^A, U^A) \models \\ \text{[LFP}_{xy, X} (\neg Ux \wedge (Exy \vee \exists z (Exz \wedge Xzy))) \\ \vee (Ux \wedge \exists z Exz \wedge \forall z (Exz \rightarrow (z = y \vee Xzy)))]xy[a, b].$$

The logic $\text{FO}(\text{ATC})$ is defined as $\text{FO}(\text{TC})$, replacing the TC-clause by

$$\frac{\varphi(\bar{x}, \bar{y}), \psi(\bar{x})}{[\text{ATC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y}), \psi(\bar{x})] \bar{s} \bar{t}}$$

where $\bar{x}, \bar{y}, \bar{s}, \bar{t}$ all have the same length. The meaning of the new formula is given by

$$(\bar{s}, \bar{t}) \in \text{ATC}(\varphi(-, -), \psi(-)).$$

We now have:

Theorem 9.4.8 $\text{FO}(\text{ATC}) \equiv \text{FO}(\text{LFP})$.

Corollary 9.4.9 $\text{FO}(\text{ATC})$ captures PTIME. □

Proof (of 9.4.8). As $\text{FO}(\text{ATC}) \leq \text{FO}(\text{LFP})$ is clear from the definition of ATC given above, we need only prove the other direction.

So let φ be an $\text{FO}(\text{LFP})$ -sentence. By 9.4.7(1) we may assume that φ is equivalent to the nontrivial sentence

$$\exists u[\text{LFP}_{\bar{x}, X} (\psi_0(\bar{x}) \vee \forall \bar{y} \exists \bar{z} \in X (\psi_1(\bar{x}, \bar{y}) \rightarrow \psi_2(\bar{y}, \bar{z})))] \bar{u},$$

where X does not appear in the quantifier-free formulas ψ_0, ψ_1, ψ_2 and where the free variables are among the displayed ones. In view of nontriviality, this formula is equivalent to

$$\exists u[\text{LFP}_{\bar{x}, X} (\psi_0(\bar{x}) \vee \forall \bar{y} (\psi_1(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \in X \psi_2(\bar{y}, \bar{z})))] \bar{u}.$$

Let F_∞ be the fixed-point of the fixed-point process in the last formula. The alternating transitive closure capturing it will make use of a binary and a unary relation on the set of all tuples $\bar{x} \bar{y} v_1 \dots v_4$. For $i = 1, \dots, 4$ we write $\bar{x} \bar{y} i$ for any such tuple $\bar{x} \bar{y} \bar{v}$ with $\epsilon_i(\bar{v})$, where the ϵ_i are as in the proof of 9.4.5. Then the relations will be defined by FO-formulas φ_E and φ_U in such

a way that $\bar{x} \in F_\infty$ iff there is an ATC-path from $\bar{x} \bar{x} 1$ to a point of the form $\bar{y} \bar{y} 4$; more precisely:

$$(*) \quad \bar{x} \in F_\infty \quad \text{iff} \quad \text{some (or equivalently, all) pairs of tuples of the form } (\bar{x} \bar{x} 1, \bar{y} \bar{y} 4) \text{ are in } \text{ATC}(\varphi_E(-, -), \varphi_U(-)).$$

As universal points we take all tuples $\bar{x} \bar{y} 1$, that is, we set

$$\varphi_U(\bar{x}, \bar{y}, \bar{v}) := \epsilon_1(\bar{v}).$$

The binary relation consists of edges that take care of the disjuncts $\psi_0(\bar{x})$ and $\forall \bar{y}(\psi_1(\bar{x}, \bar{y}) \rightarrow \exists \bar{z} \in X \psi_2(\bar{y}, \bar{z}))$ of the fixed-point formula above. Concerning $\psi_0(\bar{x})$, it contains the edges indicated by

$$\bar{x} \bar{x} 1 \xrightarrow{\psi_0(\bar{x})} \bar{x} \bar{x} 2 \longrightarrow \bar{y} \bar{y} 4,$$

and concerning the second disjunct, the edges indicated by

$$\begin{array}{ccc} & & \neg\psi_1(\bar{x}, \bar{y}) \rightarrow \bar{y}' \bar{y}' 4 \\ & & \nearrow \\ \bar{x} \bar{x} 1 \xrightarrow{\neg\psi_0(\bar{x})} & \bar{x} \bar{y} 3 & \\ & & \searrow \\ & & \psi_1(\bar{x}, \bar{y}) \rightarrow \bar{z} \bar{z} 1 \\ & & \psi_2(\bar{y}, \bar{z}) \end{array}$$

So as φ_E we take

$$\begin{aligned} \varphi_E(\bar{x} \bar{y} \bar{v}, \bar{x}' \bar{y}' \bar{v}') &:= (\epsilon_1(\bar{v}) \wedge \psi_0(\bar{x}) \wedge \bar{x} = \bar{y} = \bar{x}' = \bar{y}' \wedge \epsilon_2(\bar{v}')) \\ &\vee (\epsilon_2(\bar{v}) \wedge \bar{x} = \bar{y} \wedge \bar{x}' = \bar{y}' \wedge \epsilon_4(\bar{v})) \\ &\vee (\epsilon_1(\bar{v}) \wedge \neg\psi_0(\bar{x}) \wedge \bar{x} = \bar{y} \wedge \bar{x}' = \bar{x} \wedge \epsilon_3(\bar{v}')) \\ &\vee (\epsilon_3(\bar{v}) \wedge \neg\psi_1(\bar{x}, \bar{y}) \wedge \bar{x}' = \bar{y}' \wedge \epsilon_4(\bar{v}')) \\ &\vee (\epsilon_3(\bar{v}) \wedge \psi_1(\bar{x}, \bar{y}) \wedge \psi_2(\bar{y}, \bar{x}') \wedge \bar{x}' = \bar{y}' \wedge \epsilon_1(\bar{v}')). \end{aligned}$$

Then (*) is obvious and yields that φ is equivalent to

$$\begin{aligned} &\exists \bar{u} \exists \bar{z} \exists \bar{w} \exists \bar{w}' (\epsilon_1(\bar{w}) \wedge \epsilon_4(\bar{w}') \wedge \\ &[\text{ATC}_{\bar{x} \bar{y} \bar{v}, \bar{x}' \bar{y}' \bar{v}'} \varphi_E(\bar{x} \bar{y} \bar{v}, \bar{x}' \bar{y}' \bar{v}'), \varphi_U(\bar{x}, \bar{y}, \bar{v})] \hat{u} \hat{u} \hat{w} \bar{z} \bar{z} \bar{w}'). \end{aligned}$$

□

9.5 An Application of Negative Fixed-Point Logic

The normal form obtained in Theorem 9.4.3 tells us – as already pointed out in the remarks following this theorem – that in FO(PFP) we can restrict ourselves to PFP operations where every element of a new stage is witnessed by two elements, one belonging to the previous stage (a positive witness), the other one belonging to its complement (a negative witness). In this section we analyze the expressive power of the fragment of FO(PFP) containing PFP operators with only negative witnesses and give an application of this fragment to so-called well-founded DATALOG.

For $k, l \in \mathbb{N}$ let $\text{PFP}(k, l)$ be the class of FO(PFP)-formulas containing only fixed-point operations where every element of a new stage (besides the first one) is witnessed by k positive and l negative elements, i.e., $\varphi \in \text{FO(PFP)}$ is in $\text{PFP}(k, l)$ if all its subformulas starting with a PFP operator have the form

$$[\text{PFP}_{\bar{x}, X} (\varphi_0 \vee \exists \bar{y}_1 \in X \dots \exists \bar{y}_k \in X \exists \bar{z}_1 \notin X \dots \exists \bar{z}_l \notin X \varphi_1)] \bar{t}$$

where φ_0 and φ_1 do not contain X .

We have a complete picture of the expressive power of the different $\text{PFP}(k, l)$:

- Theorem 9.5.1** (a) For $k, l \geq 1$, $\text{PFP}(k, l) \equiv \text{PFP}(1, 1) \equiv \text{FO(PFP)}$.
 (b) For $k \geq 2$, $\text{PFP}(k, 0) \equiv \text{PFP}(2, 0) \equiv \text{FO(BFP)}$.
 (c) $\text{PFP}(1, 0) \equiv \text{FO(TC)}$.
 (d) For $l \geq 1$, $\text{PFP}(0, l) \equiv \text{PFP}(0, 1) \equiv \text{FO(LFP)}$. Moreover, every $\varphi \in \text{FO(LFP)}$ is equivalent to a totally defined formula of the form

$$\exists (\forall) u [\text{PFP}_{\bar{x}, X} \exists \bar{y} \notin X \psi(\bar{x}, \bar{y})] \bar{u}$$

where ψ is quantifier-free and does not contain X .

Proof. Part (a) holds by Theorem 9.4.3, parts (b) and (c) hold by Proposition 8.7.1. We turn to (d) and first prove $\text{PFP}(0, l) \leq \text{FO(LFP)}$ by induction on $\text{PFP}(0, l)$ -formulas. So, consider a formula $\varphi = [\text{PFP}_{\bar{x}, X} \psi] \bar{t}$ with $\psi = (\varphi_0 \vee \exists \bar{y}_1 \notin X \dots \exists \bar{y}_l \notin X \varphi_1)$ where φ_0 and φ_1 do not contain X . By induction hypothesis, we may assume that φ_0 and φ_1 are FO(LFP)-formulas. Then ψ is equivalent to the FO(LFP)-formula

$$(\varphi_0 \vee \exists \bar{y}_1 \dots \exists \bar{y}_l (\neg X \bar{y}_1 \wedge \dots \wedge \neg X \bar{y}_l \wedge \varphi_1))$$

negative in X . Therefore, by 8.3.2, $[\text{PFP}_{\bar{x}, X} \psi] \bar{t}$ is equivalent to an FO(LFP)-formula.

We sketch a proof of $\text{FO(LFP)} \leq \text{PFP}(0, 1)$ and of the normal form claimed in (d) and refer the reader to [112] for details. Let φ be a formula of FO(LFP). We may assume that

$$\varphi = \exists(\forall)u[\text{LFP}_{\bar{z},Z} \varphi_0] \tilde{u}$$

with first-order φ_0 . For simplicity, let φ be a sentence.

One can simulate the semantics of FO(LFP)-formulas by a game for a definable digraph, thereby obtaining a quantifier-free first-order formula $\psi(\bar{x}, \bar{y})$ with $\bar{x} = x_1 \dots x_k$ and $\bar{y} = y_1 \dots y_k$ for suitable k such that (1),(2), and (3) hold for all \mathcal{A} :

- (1) $(A^k, \psi^{\mathcal{A}}(-, -))$ is a digraph, i.e., $\mathcal{A} \models \forall \bar{x} \neg \psi(\bar{x}, \bar{x})$.
- (2) $\mathcal{A} \models [\text{LFP}_{\bar{z},Z} \varphi_0] \tilde{u}$ iff \tilde{u} is won in the game associated with $(A^k, \psi^{\mathcal{A}}(-, -))$ (cf. 8.1.5(b)).
- (3) $(A^k, \psi^{\mathcal{A}}(-, -))$ has no drawn points.

Therefore (cf. the example previous to 8.3.2), $\exists(\forall)u[\text{PFPP}_{\bar{x},X} \exists \bar{y} \notin X \psi(\bar{x}, \bar{y})] \tilde{u}$ is a totally defined formula equivalent to φ . □

We apply the normal form just proven to *well-founded* DATALOG, WF-DATALOG. (The exercise at the end of this section shows how one can obtain some results for FO(LFP) as immediate consequences of this normal form.) Consider the program

$$\begin{aligned} Qx &\leftarrow Qx \\ Qx &\leftarrow \neg Qx \end{aligned}$$

In the semantics of both I-DATALOG and P-DATALOG, for a structure \mathcal{A} we have $Q_\infty = A$, the elements of A getting into Q using the “wrong” information “ $\neg Q = A$ ” in the first step. In WF-DATALOG, intuitively speaking, one requires that the values obtained for the intentional symbols are “re-confirmed by their positive occurrences”. More precisely: The semantics of WF-DATALOG, the so-called *well-founded semantics*, treats positive and negative occurrences of intentional symbols in the bodies of clauses in an asymmetric way: Let Π be a WF-DATALOG program, i.e., a general logic program. For every intentional symbol P replace all occurrences of P in the heads of clauses and all positive occurrences of P in the bodies by a new relation symbol P' . Now, the original P does not further occur in the head of any clause of the resulting program Π' ; hence, P is an extensional symbol of Π' and thus, Π' is a DATALOG program. In order to simplify the presentation, we assume that Π only contains a single intentional relation symbol P .

In every τ_{ext} -structure \mathcal{A} , the program Π gives rise to a sequence $(P_n)_{n \geq 0}$ of relations on A defined by

$$\begin{aligned} P_0 &:= \emptyset, \\ P_{n+1} &\text{ is the result for } P' \text{ of the evaluation of the DATALOG program } \Pi' \text{ in } (\mathcal{A}, P_n), \text{ i.e., taking } P_n \text{ as interpretation of the extensional symbol } P \text{ of } \Pi'. \end{aligned}$$

For example, for the WF-DATALOG program

$$\begin{aligned} \Pi_0 : \quad & Pxy \leftarrow Rxyz, \neg Pux \\ & Pxy \leftarrow \neg Sxy, Rxyz, Pxz, \neg Pyy \end{aligned}$$

we have

$$\begin{aligned} \Pi'_0 : \quad & P'xy \leftarrow Rxyz, \neg Pux \\ & P'xy \leftarrow \neg Sxy, Rxyz, P'xz, \neg Pyy \end{aligned}$$

and the stages P_n of the evaluation of Π_0 as described above are just the stages F_n^φ of F^φ , where

$$\begin{aligned} \varphi(xy, P) := \quad & [\text{LFP}_{xy, P'} (\exists z \exists u (Rxyz \wedge \neg Pux) \\ & \vee \exists z (\neg Sxy \wedge Rxyz \wedge P'xz \wedge \neg Pyy))] xy. \end{aligned}$$

We come back to the general case above, i.e., to Π , the τ_{ext} -structure \mathcal{A} , and the sequence $(P_n)_{n \geq 0}$. Let P_{true} be the truth set of this sequence, i.e.,

$$P_{\text{true}} = \{a \in A \mid \exists n_0 \forall n \geq n_0 : a \in P_n\}$$

(cf. 8.3.16). Then, by definition, in the WF-semantics, Π leads to the τ -structure $\mathcal{A}(\Pi) := (\mathcal{A}, P_{\text{true}})$. For terms \bar{t} , the formula $(\Pi, P)\bar{t}$ is a WF-DATALOG formula of vocabulary τ_{ext} , its meaning being “ $\bar{t} \in P_{\text{true}}$ ”.

Note that $P_{\text{true}} = P_\infty$ if the fixed-point P_∞ of the sequence $(P_n)_{n \geq 0}$ exists. Call Π a *totally defined* WF-DATALOG program and $(\Pi, P)\bar{t}$ a *totally defined* WF-DATALOG formula, if the fixed-point exists in all structures.

In case Π is a DATALOG program, the WF-DATALOG formula $(\Pi, P)\bar{t}$ is equivalent to the DATALOG formula $(\Pi, P)\bar{t}$; in case Π is a totally WF-DATALOG program and the bodies of the clauses in Π contain only negative occurrences of P , the WF-DATALOG formula $(\Pi, P)\bar{t}$ is equivalent to the P-DATALOG formula $(\Pi, P)\bar{t}$. The same applies to the program in part (b) of the following theorem; so, in a certain sense, we can do without the well-founded semantics.

Theorem 9.5.2 (a) WF-DATALOG \equiv FO(LFP).

(b) *Every* WF-DATALOG formula is equivalent to a *totally defined* WF-DATALOG formula $(\Pi, P)\bar{x}$ with Π of the form

$$\begin{aligned} Z\bar{x}\bar{y} &\leftarrow \gamma_{11}, \dots, \gamma_{1k_1}, \neg Z\bar{x}\bar{z} \\ &\vdots \\ Z\bar{x}\bar{y} &\leftarrow \gamma_{s1}, \dots, \gamma_{sk_s}, \neg Z\bar{x}\bar{z} \\ P\bar{x} &\leftarrow Z\bar{x}\bar{u} \end{aligned}$$

where the γ_{ij} contain neither Z nor P .

Proof. (a) WF-DATALOG \leq FO(LFP): Let $(\Pi, P)\bar{t}$ be a formula of WF-DATALOG and again, for simplicity (otherwise, argue with FO(S-LFP)), assume that P is the unique intentional symbol, i.e., Π has the form

$$\begin{aligned} P\bar{x} &\leftarrow \gamma_{11}, \dots, \gamma_{1k_1} \\ &\vdots \\ P\bar{x} &\leftarrow \gamma_{s1}, \dots, \gamma_{sk_s}. \end{aligned}$$

Then the stages P_n of Π evaluated in a τ_{ext} -structure are just the stages F_n^φ of F^φ , where

$$\varphi(\bar{x}, P) := [\text{LFP}_{\bar{x}, P'} \bigvee_{i=1}^s \exists \bar{y}_i (\gamma'_{i1} \wedge \dots \wedge \gamma'_{ik_i})] \bar{x}.$$

where \bar{y}_i are the variables in $(\gamma_{i1} \wedge \dots \wedge \gamma_{ik_i})$ distinct from \bar{x} and where γ'_{ij} is obtained from γ_{ij} by replacing positive occurrences of P by P' . Then, $\varphi(\bar{x}, P)$ is negative in P . Hence, F^φ is antitone, i.e.,

$$X \subseteq Y \text{ implies } F^\varphi(Y) \subseteq F^\varphi(X)$$

(see Exercise 8.1.6). Therefore,

$$F_0^\varphi \subseteq F_2^\varphi \subseteq F_4^\varphi \subseteq \dots \subseteq F_5^\varphi \subseteq F_3^\varphi \subseteq F_1^\varphi$$

(cf. 8.1.5(a)). Clearly, $G := F^\varphi \circ F^\varphi$ is monotone, and $G = F^\psi$ holds for the formula $\psi(\bar{x}, P) := \varphi(\bar{x}, \varphi(-, P))$ positive in P (cf. 8.3.2). Since $F_n^\psi = F_{2n}^\varphi$, we have $F_\infty^\psi = \bigcup_{n \geq 0} F_{2n}^\varphi = P_{\text{true}}$. Hence, $(\Pi, P)\bar{t}$ is equivalent to the FO(LFP)-formula

$$[\text{LFP}_{\bar{x}, P} \psi(\bar{x}, P)] \bar{t}.$$

FO(LFP) \leq WF-DATALOG: Let $\varphi(\bar{x})$ be an FO(LFP)-formula. By Theorem 9.5.1(d), it is equivalent to a totally defined FO(PFP)-formula of the form

$$\exists u [\text{PFP}_{\bar{y}, X} \exists \bar{z} \notin X \psi(\bar{x}, \bar{y}, \bar{z})] \tilde{u}$$

where ψ is quantifier-free and does not contain X . Then, $\varphi(\bar{x})$ is equivalent to

$$\exists u [\text{PFP}_{\bar{x}, \bar{y}, Z} \exists \bar{z} (\neg Z \bar{x} \bar{z} \wedge \psi(\bar{x}, \bar{y}, \bar{z}))] \bar{x} \tilde{u}.$$

We can assume that ψ is in disjunctive normal form, e.g., $\psi = \bigvee_{i=1}^k (\gamma_{i1} \wedge \dots \wedge \gamma_{ik_i})$ with atomic or negated atomic γ_{ij} . Then, $\varphi(\bar{x})$ is equivalent to

$$\exists u [\text{PFP}_{\bar{x}, \bar{y}, Z} \bigvee_{i=1}^k \exists \bar{z} (\neg Z \bar{x} \bar{z} \wedge \gamma_{i1} \wedge \dots \wedge \gamma_{ik_i})] \bar{x} \tilde{u}$$

and hence, equivalent to the totally defined WF-DATALOG formula $(\Pi, P)\bar{x}$, where Π consists of the clauses

$$\begin{aligned}
 Z\bar{x}\bar{y} &\leftarrow \gamma_{11}, \dots, \gamma_{1k_1}, \neg Z\bar{x}\bar{z} \\
 &\vdots \\
 &\vdots \\
 Z\bar{x}\bar{y} &\leftarrow \gamma_{s1}, \dots, \gamma_{sk_s}, \neg Z\bar{x}\bar{z} \\
 P\bar{x} &\leftarrow Z\bar{x}\bar{u}
 \end{aligned}$$

This completes the proof of part (a). At the same time, we have proved part (b), too. \square

The following exercise contains two further applications of part (d) of Theorem 9.5.1: In part (a) we give a short proof for the normal form for FO(LFP) contained in Theorem 9.4.7 (thereby restricting ourselves to (2)); in part (b) we show that every FO(LFP) formula is equivalent to a first-order formula containing relations implicitly definable in FO (cf. 8.5.8).

Exercise 9.5.3 Let $\rho(\bar{w})$ be an FO(LFP)-formula. According to 9.5.1(d) choose a totally defined formula equivalent to $\rho(\bar{w})$ of the form

$$(*) \quad \exists(\forall)u[\text{PFP}_{\bar{x},X}\exists\bar{y} \notin X\chi] \bar{u}$$

where χ is quantifier-free, does not contain X , and $\text{free}(\chi) \subseteq \{\bar{x}, \bar{y}, \bar{w}\}$.

(a) Set $\varphi(\bar{x}, X) := \exists\bar{y} \notin X\chi(\bar{x}, \bar{y})$ and $\psi(\bar{x}, X) := \varphi(\bar{x}, \varphi(-, X))$. Then, $F^\psi = F^\varphi \circ F^\varphi$. Since the formula in (*) is totally defined, we have $F_\infty^\psi = F_\infty^\varphi$ (cf. 8.1.5 and 8.3.2). Show that $\psi(\bar{x}, X)$ is equivalent to

$$\exists\bar{y}(\chi(\bar{x}, \bar{y}) \wedge \forall\bar{z}(\chi(\bar{y}, \bar{z}) \rightarrow X\bar{z})).$$

Conclude that

$$\models \exists\bar{z}\neg X\bar{z} \rightarrow (\psi(\bar{x}, X) \leftrightarrow \exists\bar{y}\forall\bar{z} \notin X(\chi(\bar{x}, \bar{y}) \wedge \neg\chi(\bar{y}, \bar{z})))$$

and hence, that $\rho(\bar{w})$ is equivalent

$$\exists(\forall)u[\text{LFP}_{\bar{x},X}\exists\bar{y}\forall\bar{z} \notin X(\chi(\bar{x}, \bar{y}) \wedge \neg\chi(\bar{y}, \bar{z}))] \bar{u}.$$

(b) Clearly, $\rho(\bar{w})$ is equivalent to to the formula

$$\exists(\forall)\bar{u}[\text{PFP}_{\bar{x}\bar{w},Y}\exists\bar{y}(\neg Y\bar{y}\bar{w} \wedge \chi)] \bar{u}\bar{w},$$

that is totally defined, too. Introduce $\varphi(\bar{x}\bar{w}, Y)$ as above. By 8.1.5, F_∞^φ is the unique fixed-point of F^φ . Let $\psi_0(Y) := \forall\bar{x}\bar{w}(Y\bar{x}\bar{w} \leftrightarrow \exists\bar{y}(\neg Y\bar{y}\bar{w} \wedge \chi))$ and conclude that

$$\models_{\text{fin}} \exists^{-1}Y\psi_0(Y) \quad \text{and} \quad \models_{\text{fin}} \forall Y(\psi_0(Y) \rightarrow (\rho(\bar{w}) \leftrightarrow Y\bar{w}_1\bar{w})),$$

i.e., that $\rho(\bar{w})$ is equivalent to a section of a relation implicitly definable in FO. \square

9.6 Hierarchies of Fixed-Point Logics

We first discuss the problem to what extent the expressive power of fixed-point operators depends on their arities. This problem comes up very naturally with the proofs of the normal form theorems, where the simplifications we have obtained with respect to the number of fixed-point operators, quantifiers, and occurrences of second-order variables had to be paid by enlarging the arity of the fixed-point operators involved.

Denote by $\text{FO}(\text{LFP}^r)$, $\text{FO}(\text{IFP}^r)$, and $\text{FO}(\text{PFP}^r)$ the set of formulas of the corresponding fixed-point logic that contain second-order variables only of arity $= r$.⁴ And let $\text{FO}(\text{DTC}^r)$ be the fragment of $\text{FO}(\text{DTC})$ consisting of those formulas, where for any subformula of the form $[\text{DTC}_{\bar{x}, \bar{y}} \varphi] \bar{s} \bar{t}$ we have $\text{length}(\bar{x}) = r$. Define $\text{FO}(\text{TC}^r)$ similarly. Clearly,

$$(*) \quad \text{FO}(\text{DTC}^r) \leq \text{FO}(\text{TC}^r) \leq \text{FO}(\text{LFP}^r) \leq \text{FO}(\text{IFP}^r) \leq \text{FO}(\text{PFP}^r).$$

Theorem 9.6.1 *Let \mathcal{L} be one of the logics $\text{FO}(\text{DTC})$, $\text{FO}(\text{TC})$, $\text{FO}(\text{LFP})$, $\text{FO}(\text{IFP})$, $\text{FO}(\text{PFP})$. For $r \geq 1$, let \mathcal{L}^r be the corresponding fragment satisfying the arity restriction above. Then*

$$\mathcal{L}^0 < \mathcal{L}^1 < \mathcal{L}^2 < \dots$$

More precisely: If τ contains an at least binary relation symbol then

$$\mathcal{L}^0[\tau] < \mathcal{L}^1[\tau] < \dots$$

Using dummy variables, it is clear that $\mathcal{L}^0 \leq \mathcal{L}^1 \leq \mathcal{L}^2 \leq \dots$. In view of (*), the strict inequalities are an immediate consequence of

Theorem 9.6.2 *For any $r \geq 0$ we have*

$$\text{FO}(\text{TC}^{r+1}) \not\leq \text{FO}(\text{PFP}^r) \quad \text{and} \quad \text{FO}(\text{DTC}^{r+1}) \not\leq \text{FO}(\text{TC}^r)$$

on the class of graphs.

For a proof we refer to [64]. We only sketch the main idea. To obtain $\text{FO}(\text{TC}^{r+1}) \not\leq \text{FO}(\text{PFP}^r)$, one considers graphs definable on the set of $(r+1)$ -tuples of the universe of certain structures. These structures are sufficiently homogeneous with respect to r -tuples to cause a collapse of the expressive power of $\text{FO}(\text{PFP}^r)$ to FO . As a consequence, the query expressing that two $(r+1)$ -tuples are connected by a path is not definable in $\text{FO}(\text{PFP}^r)$ but, of course, it is expressible in $\text{FO}(\text{TC}^{r+1})$. The construction is quite intricate. It still has to be refined to obtain structures showing $\text{FO}(\text{DTC}^{r+1}) \not\leq \text{FO}(\text{TC}^r)$. Note that the first claim cannot be strengthened to $\text{FO}(\text{DTC}^{r+1}) \not\leq \text{FO}(\text{PFP}^r)$, as $\text{FO}(\text{DTC}) \leq \text{FO}(\text{PFP}^2)$ by the next proposition. \square

⁴ Thus $\text{FO}(\text{LFP}^1)$ is monadic fixed-point logic $\text{FO}(\text{M-LFP})$ as introduced in section 8.5.

Throughout this section, given a formula $\varphi(\bar{x}, \bar{y})$, we abbreviate by $\varphi_{\text{det}}(\bar{x}, \bar{y})$ the formula

$$\varphi_{\text{det}}(\bar{x}, \bar{y}) := \varphi(\bar{x}, \bar{y}) \wedge \forall \bar{z}(\varphi(\bar{x}, \bar{z}) \rightarrow \bar{z} = \bar{y})$$

(where \bar{z} are new variables). Then

$$\models_{\text{fin}} [\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{x} \bar{y} \leftrightarrow [\text{TC}_{\bar{x}, \bar{y}} \varphi_{\text{det}}(\bar{x}, \bar{y})] \bar{x} \bar{y}.$$

Proposition 9.6.3 $\text{FO}(\text{DTC}) \leq \text{FO}(\text{PFP}^2)$.

Proof. For r -tuples \bar{a}, \bar{e} let

$$X_{\bar{e}}(\bar{a}) := \{(a_1, e_1), \dots, (a_r, e_r)\}.$$

As the main step of an inductive proof we show how to express a DTC operator by a PFP^2 operator: Given an $\text{FO}(\text{PFP}^2)$ -formula $\varphi(\bar{x}, \bar{y})$, we shall find an $\text{FO}(\text{PFP}^2)$ -formula $\psi(x, y, X, \bar{x}, \bar{y}, \bar{z})$ with binary X and with parameters $\bar{x}, \bar{y}, \bar{z}$ such that for distinct elements z_1, \dots, z_r ,

(1) if there is a deterministic φ -path from \bar{x} to \bar{y} , say $\bar{x} = \bar{x}_0, \dots, \bar{x}_l = \bar{y}$, then

$F_0^\psi, F_1^\psi, \dots$ is the sequence $\emptyset, X_{\bar{z}}(\bar{x}_1), X_{\bar{z}}(\bar{x}_2), \dots, X_{\bar{z}}(\bar{x}_l), X_{\bar{z}}(\bar{x}_l), \dots$;

(2) if there is no deterministic φ -path from \bar{x} to \bar{y} then $F_\infty^\psi = \emptyset$.

Then, in structures with at least r elements (the others can be treated separately), $[\text{DTC}_{\bar{x}, \bar{y}} \varphi] \bar{x} \bar{y}$ and

$$\exists z_1 \dots \exists z_r \left(\bigwedge_{1 \leq i < j \leq r} \neg z_i = z_j \wedge \exists u \exists v [\text{PFP}_{x, y, X} \psi(x, y, X, \bar{x}, \bar{y}, \bar{z})] uv \right)$$

are equivalent, and we are done.

As $\psi(x, y, X, \bar{x}, \bar{y}, \bar{z})$ we can take the formula (note the case of a φ -path

$$\bar{x} = \bar{x}_0, \bar{x}_1, \dots, \bar{x}_m, \bar{x}_m, \dots$$

not containing \bar{y} needs special care)

$$\begin{aligned} & (X = \emptyset \wedge \exists \bar{w}(\varphi_{\text{det}}(\bar{x}, \bar{w}) \wedge (\neg \bar{w} = \bar{x} \vee \bar{w} = \bar{y})) \wedge \bigvee_{1 \leq i \leq r} (x = w_i \wedge y = z_i)) \\ & \vee \exists \bar{v} \exists \bar{w} (\bigwedge_{1 \leq i \leq r} X v_i z_i \wedge \varphi_{\text{det}}(\bar{v}, \bar{w}) \wedge \neg \bar{v} = \bar{y} \wedge \neg \bar{v} = \bar{w} \wedge \\ & \quad \bigvee_{1 \leq i \leq r} (x = w_i \wedge y = z_i)) \\ & \vee (\bigwedge_{1 \leq i \leq r} X y_i z_i \wedge \bigvee_{1 \leq i \leq r} (x = y_i \wedge y = z_i)). \end{aligned}$$

□

We now turn to the hierarchy problem for *ordered* structures. We present two propositions which will help us to discuss the general problem.

Proposition 9.6.4 *On ordered structures, $\text{FO}(\text{DTC}) \leq \text{FO}(\text{LFP}^3)$.*⁵

Proof. Consider a formula $\varphi(\bar{x}, \bar{y})$, where $\bar{x} = x_1 \dots x_r$ and $\bar{y} = y_1 \dots y_r$. In a structure \mathcal{A} of cardinality n , the (shortest) deterministic φ -path between two tuples must have length $\leq n^r$. Assume $A = \{1, \dots, n\}$. First suppose that there is a deterministic φ -path

$$(1) \quad \bar{x} = \bar{x}_0, \bar{x}_1, \dots, \bar{x}_k = \bar{y}$$

from \bar{x} to \bar{y} with $k \leq n$ (where $\bar{x}_i = x_{i1} \dots x_{ir}$). To rewrite φ as an $\text{FO}(\text{LFP}^3)$ -formula, we shall code the existence of such a φ -path by an LFP operator with stages

$$\begin{aligned} X_1 &= \{(x_{11}, 1, 1), (x_{12}, 1, 2), \dots, (x_{1r}, 1, r)\} \\ X_2 &= X_1 \cup \{(x_{21}, 2, 1), (x_{22}, 2, 2), \dots, (x_{2r}, 2, r)\} \\ &\vdots \\ X_k &= X_{k-1} \cup \{(x_{k1}, k, 1), (x_{k2}, k, 2), \dots, (x_{kr}, k, r)\} \end{aligned}$$

and $X_k = X_{k+1} = \dots$. Using s nested LFP operators appropriately, we can express the existence of a deterministic φ -path of length $\leq n^s$. For $s = r$, we get the desired formula.

To give the details of the elimination of a DTC operator by LFP^3 operators (as the main step of an inductive proof), let $\varphi(\bar{x}, \bar{y})$ be any $\text{FO}(\text{LFP}^3)$ formula.

Consider the formula

$$\psi(\bar{x}, \bar{y}) := [\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{x} \bar{y}.$$

Then, in structures with at least $\max\{3, r\}$ elements, the statement

“there is a deterministic φ -path of length \leq (cardinality of the universe) from \bar{x} to \bar{y} ”

can be expressed by the $\text{FO}(\text{LFP}^3)$ -formula $\varphi^*(\bar{x}, \bar{y})$ below, which uses the above coding of the path (1) prolonged to the path $\bar{x} = \bar{x}_0, \bar{x}_1, \dots, \bar{x}_k = \bar{y}, \bar{y}, \dots, \bar{y}$ of length n :⁶

$$\begin{aligned} \varphi^* &:= \exists x \exists v [\text{LFP}_{xuv, X} \\ &\quad \exists \bar{y}' (\varphi_{\text{det}}(\bar{x}, \bar{y}') \wedge u = \min + 1 \wedge \bigvee_{1 \leq i \leq r} (x = y'_i \wedge v = i)) \\ &\quad \vee \exists u' \exists \bar{x}' \exists \bar{y}' (\varphi_{\text{det}}(\bar{x}', \bar{y}') \wedge \neg \bar{x}' = \bar{y} \wedge \bigwedge_{1 \leq i \leq r} X x'_i u' i \wedge u = u' + 1 \\ &\quad \wedge \bigvee_{1 \leq i \leq r} (x = y'_i \wedge v = i) \wedge (u = \max \rightarrow \bar{y}' = \bar{y})) \\ &\quad \vee \exists u' (\bigwedge_{1 \leq i \leq r} X y_i u' i \wedge u = u' + 1 \wedge \bigvee_{1 \leq i \leq r} (x = y_i \wedge v = i))] x \max v. \end{aligned}$$

⁵ The next exercise shows that LFP^3 can be replaced by LFP^2 .

⁶ As usual, for $i \geq 1$ we abbreviate by $z = i$ that z is the i -th element in the ordering.

We now iterate the transition from φ to φ^* , defining

$$\varphi^1 = \varphi^* \quad \text{and} \quad \varphi^{n+1} = (\varphi^n)^*.$$

Then $\varphi^s(\bar{x}, \bar{y})$ expresses

“there is a deterministic φ -path of length \leq (cardinality of the universe)^s from \bar{x} to \bar{y} .”

Hence, φ^r is an FO(LFP³)-formula equivalent to $[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{x} \bar{y}$. \square

Exercise 9.6.5 Show that $\text{FO}(\text{DTC}) \leq \text{FO}(\text{LFP}^2)$ on ordered structures. Hint: Using arithmetics definable in $\text{FO}(\text{LFP}^2)$, code the last two components of X in the preceding proof by a single one. \square

We know that $\text{FO}(\text{DTC}) \not\leq \text{FO}(\text{LFP}^1)$ on ordered structures, since even $\text{FO}(\text{DTC}) \not\leq \text{MSO}$ on ordered structures (cf. 8.6.3). But we have:

Proposition 9.6.6 *On ordered structures, $\text{FO}(\text{DTC}) \leq \text{FO}(\text{PFP}^1)$.*

Proof. Consider the formula

$$[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}$$

with $\bar{x} = x_1 \dots x_r$ and $\bar{y} = y_1 \dots y_r$. We show how the DTC operator can be expressed by PFP^1 . For this purpose we consider a deterministic φ -path

$$(*) \quad \bar{s} = \bar{x}_0, \bar{x}_1, \bar{x}_2, \dots$$

starting from \bar{s} . The idea is to code r -tuples \bar{x} by subsets $P_{\bar{x}}$ of the universe and to find a formula of $\text{FO}(\text{PFP}^1)$ whose PFP^1 operator has the stages $\emptyset, P_{\bar{x}_1}, P_{\bar{x}_2}, \dots, P_{\bar{x}_k}, P_{\bar{t}}, P_{\bar{t}}, \dots$ in case $\bar{x}_k = \bar{t}$ in (*), but does not reach a fixed-point in case the deterministic φ -path started in \bar{s} does not pass \bar{t} .

We present the codification of \bar{x} in a set $P_{\bar{x}}$: Let (the interpretations of) \bar{x} be given in an ordered structure. For $1 \leq i \leq r$, let m_i be the position of x_i in the ordering induced on the x_i (for example, if $r = 4$ and $x_3 < x_1 = x_4 < x_2$ then $m_1 = 2, m_2 = 3, m_3 = 1, m_4 = 2$). Note that $1 \leq m_i \leq r$. Order the set $\{(n_1, \dots, n_r) \mid 1 \leq n_1, \dots, n_r \leq r\}$, a set independent of the given structure, lexicographically. If (m_1, \dots, m_r) is the l -th element in this ordering set

$$P_{\bar{x}} := \{\bar{x}\} \cup I$$

where I is the leftmost interval in the given structure having length $r + 1 + l$ and containing no x_i (such an interval exists in structures of cardinality $\geq (r+1) \cdot (r+2+r^r)$). Note that $P_{\bar{x}}$ uniquely codes $\{\bar{x}\}$ and l and hence, \bar{x} . Since the coding and decoding of \bar{x} in $P_{\bar{x}}$ is first-order definable, it is easy to write down an $\text{FO}(\text{PFP}^1)$ -formula equivalent to $[\text{DTC}_{\bar{x}, \bar{y}} \varphi(\bar{x}, \bar{y})] \bar{s} \bar{t}$. \square

The hierarchy result for $\text{FO}(\text{PFP})$ remains true also on ordered structures.

Theorem 9.6.7 *For all $r \geq 0$, $\text{FO}(\text{PFP}^r) < \text{FO}(\text{PFP}^{r+1})$ on ordered structures.*

Proof. For $e \in \mathbb{R}, e \geq 1$, consider the function $n \mapsto n^e$. We say that a class K is in $\text{SPACE}(n^e)$ if for some $c > 0$ there is an $c \cdot n^e$ space-bounded deterministic Turing machine accepting K . From complexity theory we know that for $r \in \mathbb{N}$,

$$(*) \quad \text{SPACE}(n^r) \neq \text{SPACE}(n^{r+\frac{1}{2}}).$$

We show that

$$(1) \quad \text{each class axiomatizable in } \text{FO}(\text{PFP}^r) \text{ is in } \text{SPACE}(n^r)$$

and for $\epsilon \in \mathbb{R}, 0 \leq \epsilon < 1$ that

$$(2) \quad \text{each class in } \text{SPACE}(n^{r+\epsilon}) \text{ is axiomatizable in } \text{FO}(\text{PFP}^{r+1}).$$

Then we are finished: If K is a class in $\text{SPACE}(n^{r+\frac{1}{2}}) \setminus \text{SPACE}(n^r)$ then K is axiomatizable in $\text{FO}(\text{PFP}^{r+1})$ (by (2)), but not in $\text{FO}(\text{PFP}^r)$ (by (1)).

Claim (1) is already inherent in the proof of 7.4.2. Claim (2) is obtained by an analysis of the proof showing that classes of ordered structures in PSPACE are axiomatizable in $\text{FO}(\text{PFP})$ (cf. 7.3.2). We sketch a proof. To code the configurations of an n^r space-bounded machine, in section 7.3 we used a $(2+r)$ -ary relation, where the first two components served to distinguish the distinct ingredients of configurations, their number k being independent of the input structure. Clearly, k distinct numbers of a single component are sufficient for this purpose. The remaining $n-k$ numbers of this component can be used (in sufficiently large structures) to encode the n^ϵ part of an $n^{r+\epsilon}$ space-bounded machine (because of the preceding proposition and 7.3.11 the arithmetics needed for the encoding is available in $\text{FO}(\text{PFP}^{r+1})$). \square

Corollary 9.6.8 *On ordered structures, $\text{FO}(\text{DTC}) < \text{FO}(\text{PFP})$.*

Proof. On ordered structures, we have $\text{FO}(\text{DTC}) \leq \text{FO}(\text{PFP}^1)$ (cf. 9.6.6) and $\text{FO}(\text{PFP}^1) < \text{FO}(\text{PFP})$ by the preceding theorem. \square

In the following remarks we restrict ourselves to ordered structures. It is open whether the arity hierarchy is strict for $\text{FO}(\text{DTC})$, $\text{FO}(\text{TC})$, or $\text{FO}(\text{LFP})$, that is, whether the analogues of 9.6.7 are true for these logics. The problem for $\text{FO}(\text{LFP})$ is related to prominent questions in complexity theory: If $\text{FO}(\text{LFP}) \equiv \text{FO}(\text{LFP}^r)$ for some $r \geq 1$, then by 9.6.7, $\text{FO}(\text{LFP}) \neq \text{FO}(\text{PFP})$ and hence, $\text{PTIME} \neq \text{PSPACE}$ (cf. 7.5.2). On the other hand, in view of $\text{FO}(\text{DTC}) \leq \text{FO}(\text{LFP}^3)$ (cf. 9.6.4), if $\text{FO}(\text{LFP}) \neq \text{FO}(\text{LFP}^3)$ (by 9.6.5, even if $\text{FO}(\text{LFP}) \neq \text{FO}(\text{LFP}^2)$) then $\text{FO}(\text{DTC}) \neq \text{FO}(\text{LFP})$ and hence, $\text{LOGSPACE} \neq \text{PTIME}$ (cf. 7.5.2).

Exercise 9.6.9 Let $\mathcal{L} \in \{\text{FO}(\text{DTC}), \text{FO}(\text{TC}), \text{FO}(\text{LFP})\}$. Show $\mathcal{L}^1 < \mathcal{L}^2$ on ordered structures (cf. 8.6.3). Show that $\text{FO}(\text{LFP}^1) < \text{FO}(\text{IFP}^1)$. Hint: Use 6.3.3 (and 6.2.3) to show that $\text{FO}(\text{IFP}^1) \not\leq \text{MSO}$ on word models. \square

Exercise 9.6.10 Show $\text{FO}(\text{DTC}) \leq \text{FO}(\text{IFP}^2)$. Hint: By 9.6.5, $\text{FO}(\text{DTC}) \leq \text{FO}(\text{LFP}^2)$ on ordered structures. Use deterministic paths to define sufficiently large orderings that allow to carry out the corresponding argument. \square

We know that every $\text{FO}(\text{LFP})$ -formula is equivalent to an $\text{FO}(\text{LFP})$ -formula containing just one fixed-point operation. But, given r , does this result hold for $\text{FO}(\text{LFP}^r)$ instead of $\text{FO}(\text{LFP})$? In other words: Do we really have to increase the arity when replacing nested fixed-point operations by single ones? In general, we have. We close by stating a theorem which contains the corresponding result for all fixed-point logics. For a proof we refer to [113].

For an $\text{FO}(\text{DTC}^1)$ -formula φ let $n(\varphi)$ be the maximal number of nested DTC operations in φ . Define $\rho : \text{FO}(\text{PFP}) \rightarrow \mathbb{N}$ by

$$\begin{aligned} \rho(\varphi) &:= 0, & \text{if } \varphi \text{ is atomic;} \\ \rho(\neg\varphi) &:= \rho(\varphi); \\ \rho(\varphi \vee \psi) &:= \max\{\rho(\varphi), \rho(\psi)\}; \\ \rho(\exists x\varphi) &:= \rho(\varphi); \\ \rho([\text{PFP}_{\bar{x}, X}\varphi]\bar{t}) &:= \rho(\varphi) + \text{arity of } X. \end{aligned}$$

Theorem 9.6.11 For $k \geq 0$ there is an $\text{FO}(\text{DTC}^1)$ -sentence φ with $n(\varphi) = k + 1$, which is not equivalent to any $\text{FO}(\text{PFP})$ -sentence ψ with $\rho(\psi) \leq k$. \square

Corollary 9.6.12 For every $k \geq 1$ there is an $\text{FO}(\text{LFP}^1)$ -sentence not equivalent to any $\text{FO}(\text{LFP}^k)$ -sentence with a single LFP operation. \square

Notes 9.6.13 A reference for databases and logic programs is [1]. The main reference for sections 9.1 and 9.2 is [3]. Theorem 9.3.2 goes back to [7]. Most of the results of section 9.4 are due to Grohe and contained in his thesis [62]; cf. also [63]. References for section 9.5 are [112] (for Theorem 9.5.1(d)) and [42] (for Theorem 9.5.2). The main result of section 9.6, Theorem 9.6.2, is due to Grohe [62, 64]. Further references are [67, 88].

10. Optimization Problems

Many of the oldest and more prominent examples of NP-TIME-complete decision problems arose from the study of combinatorial optimization problems, the NP-TIME-completeness reflecting their apparent intractability. More precisely, the NP-TIME-completeness of the decision problem rules out the existence of a polynomial time algorithm for the optimization problem (unless P-TIME = NP-TIME). Of course, the intractability of an optimization problem does not exclude the existence of efficient algorithms that provide approximative solutions. And in fact, such algorithms may be necessary for practical purposes.

It has turned out that intractable optimization problems may behave quite differently with respect to the existence of approximative algorithms (a phenomenon that led to various notions of approximability).

In this chapter we first show that the class of polynomially bounded optimization problems coincides with the class of optimization problems definable by first-order formulas. This characterization leads to a logical or descriptive classification of those problems. Section 10.2 presents some first results documenting a relationship between this classification and approximability. For further study we refer the reader to the literature.

10.1 Polynomially Bounded Optimization Problems

To prepare our definition of an optimization problem, let us consider two examples, a maximization and a minimization problem.

A *cut* C in a graph $\mathcal{G} = (G, E^G)$ is a subset of G .

MAXCUT: Given a graph \mathcal{G} , MAXCUT asks for the maximal number of edges between the two parts of a cut, that is, for the maximum of

$$|\{(a, b) \mid E^G ab, a \in C, b \notin C\}|,$$

where C ranges over all cuts in \mathcal{G} .

A *vertex cover* C in a graph $\mathcal{G} = (G, E^G)$ is a subset of G such that

$$\mathcal{G} \models \forall x \forall y (Exy \rightarrow (Cx \vee Cy)).$$

MINVERTEX: Given a graph \mathcal{G} , MINVERTEX asks for the minimum of $\|C\|$ where C ranges over all vertex covers in \mathcal{G} .

In this chapter we use some sloppy formulations. For example, we say that a class K of (unordered) structures is acceptable in polynomial time, actually meaning that the class $K_{<}$ of its ordered representations (cf. 7.5.11) is in PTIME. And if we say that a function $f : K \rightarrow \mathbb{N}$ is computable in polynomial time, we mean that the function $f_{<} : K_{<} \rightarrow \mathbb{N}$ with

$$f_{<}((\mathcal{A}, <)) := f(\mathcal{A})$$

is computable in polynomial time.

Definition 10.1.1 An NPTIME *optimization problem* \mathcal{Q} is given by the data K, F, cost, μ , by short: $\mathcal{Q} = (K, F, \text{cost}, \mu)$, where

- K is a class of structures (of a fixed vocabulary) acceptable in polynomial time, the class of *input structures* or *input instances*.
- F is a function defined on K ; for every instance \mathcal{A} , $F(\mathcal{A})$ is the set of *feasible solutions* for \mathcal{A} . There is an $s \geq 1$ such that for every $\mathcal{A} \in K$ and $S \in F(\mathcal{A})$,

$$S \subseteq A^s.$$

Moreover, the class

$$\{(\mathcal{A}, S) \mid \mathcal{A} \in K, S \in F(\mathcal{A})\}$$

is acceptable in polynomial time.

- cost is a polynomial time computable function defined on the class $\{(\mathcal{A}, S) \mid \mathcal{A} \in K, S \in F(\mathcal{A})\}$. The values of cost are natural numbers.
- $\mu \in \{\max, \min\}$.

If $\mu = \max$ ($\mu = \min$) we speak of a *maximization* (*minimization*) problem. □

For \mathcal{Q} we define the function $\text{opt}_{\mathcal{Q}}$ on the class K of input instances by

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) := \mu\{\text{cost}(\mathcal{A}, S) \mid S \in F(\mathcal{A})\}.$$

An NPTIME optimization problem \mathcal{Q} induces the following decision problem: Given an input instance \mathcal{A} and a $k \geq 0$, is there a feasible solution S for \mathcal{A} such that $\text{cost}(\mathcal{A}, S) \geq k$ if $\mu = \max$, and $\text{cost}(\mathcal{A}, S) \leq k$ if $\mu = \min$? The name “NPTIME optimization problem” is justified by the observation that this decision problem is in NPTIME.

Let us see how the above examples, MAXCUT and MINVERTEX, fit into these definitions. In both examples the class of instances is the class of graphs. For MAXCUT the set of feasible solutions $F(\mathcal{G})$ is the power set $\text{Pow}(G)$, and the cost function is given by

$$\text{cost}(\mathcal{G}, C) = \|\{(a, b) \in E^G \mid a \in C \text{ and } b \notin C\}\|$$

and $\mu = \max$. For MINVERTEX, $F(\mathcal{G})$ is the set of vertex covers in \mathcal{G} ,

$$\text{cost}(\mathcal{G}, C) = \|C\|$$

and $\mu = \min$. In both cases the problem asks for $\text{opt}_{\mathcal{Q}}(\mathcal{G})$.

An NPTIME optimization problem \mathcal{Q} is said to be *polynomially bounded* if there is a polynomial $p \in \mathbb{N}[x]$ such that $\text{opt}_{\mathcal{Q}}(\mathcal{A}) \leq p(\|\mathcal{A}\|)$ for all input structures \mathcal{A} .

For polynomially bounded optimization problems the bridge to logic is given by the following concept.

Definition 10.1.2 An NPTIME optimization problem \mathcal{Q} is *first-order definable* if there is a first-order formula $\varphi(\bar{x}, \bar{Y})$ with free variables among $\bar{x} = x_1 \dots x_k$ and $\bar{Y} = Y_1 \dots Y_m$ such that for any input structure \mathcal{A} ,

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \mu_{\bar{R}} \|\{\bar{a} \in A \mid \mathcal{A} \models \varphi(\bar{a}, \bar{R})\}\|. \quad \square$$

Clearly, MAXCUT is first-order definable by the formula

$$\varphi(x, y, Y) := Exy \wedge Yx \wedge \neg Yy$$

and similarly, MINVERTEX by

$$\varphi(x, Y) := \forall x \forall y (Exy \rightarrow (Yx \vee Yy)) \rightarrow Yx.$$

Theorem 10.1.3 *An NPTIME optimization problem is polynomially bounded iff it is first-order definable.*

Proof. Clearly, if an optimization problem \mathcal{Q} is first-order definable, say by $\varphi(x_1, \dots, x_k, \bar{Y})$, then $\text{opt}_{\mathcal{Q}}(\mathcal{A}) \leq \|\mathcal{A}\|^k$ for all instances \mathcal{A} ; hence, \mathcal{Q} is polynomially bounded.

Now suppose that \mathcal{Q} is a polynomially bounded optimization problem, say, $\text{opt}_{\mathcal{Q}}(\mathcal{A}) \leq \|\mathcal{A}\|^k$ for all instances \mathcal{A} . Denote by τ the vocabulary of the input structures of \mathcal{Q} and let P be a new k -ary relation symbol. Set

$$K_1 := \{(\mathcal{A}, P^A) \mid \mathcal{A} \in K, P^A \subseteq A^k, \text{ there is a feasible solution } S \in F(\mathcal{A}) \text{ such that } \text{cost}(\mathcal{A}, S) \geq \|P^A\|\}$$

if $\mu = \max$, and

$$K_1 := \{(\mathcal{A}, P^A) \mid \mathcal{A} \in K, P^A \subseteq A^k, \text{ there is a feasible solution } S \in F(\mathcal{A}) \text{ such that } \text{cost}(\mathcal{A}, S) \leq \|P^A\|\}$$

if $\mu = \min$. By the remarks following Definition 10.1.1 we know that the class K_1 is in NPTIME. Therefore, by 7.5.14, there is a Σ_1^1 -formula $\exists \bar{X} \psi(\bar{X}, P)$ of vocabulary $\tau \cup \{P\}$ with first-order ψ such that for any structure (\mathcal{A}, P^A) ,

$$(\mathcal{A}, P^A) \in K_1 \text{ iff } (\mathcal{A}, P^A) \models \exists \bar{X} \psi(\bar{X}, P).$$

Thus,

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \mu\{\|P\| \mid \bar{R}, P \text{ relations over } A \text{ such that } \mathcal{A} \models \psi(\bar{R}, P)\}.$$

Therefore, in case $\text{opt} = \max$ we have

$$(1) \quad \text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_{\bar{R}, P} \|\{\bar{a} \in A \mid \mathcal{A} \models P\bar{a} \wedge \psi(\bar{R}, P)\}\|,$$

and in case $\text{opt} = \min$,

$$(2) \quad \text{opt}_{\mathcal{Q}}(\mathcal{A}) = \min_{\bar{R}, P} \|\{\bar{a} \in A \mid \mathcal{A} \models \psi(\bar{R}, P) \rightarrow P\bar{a}\}\|.$$

Hence, \mathcal{Q} is first-order definable. □

Denote by MAX PB the class of polynomially bounded maximization problems and by MAX Σ_n the class of first-order definable maximization problems, where the defining formula (cf. 10.1.2) can be chosen Σ_n . Classes such as MAX Π_n , MIN PB, ... are defined similarly (for the definition of Σ_n and Π_n compare 1.B).

Using so-called Skolem relations, one easily shows that every Σ_1^1 -formula is logically equivalent to a formula of the form $\exists \bar{X} \psi$ with $\psi \in \Pi_2$. E.g., for $\psi = \forall \bar{y} \exists z \forall \bar{u} \exists v \varphi$ with quantifier-free φ , we show how the formula $\exists \bar{X} \psi$ can be transformed into the desired form, thereby exemplifying the main step of an inductive proof. For a new relation variable Z , the formula $\exists \bar{X} \psi$ is equivalent to $\exists \bar{X} \exists Z \forall \bar{y} (\exists w Z \bar{y} w \wedge \forall z (Z \bar{y} z \rightarrow \forall \bar{u} \exists v \varphi))$ and hence, equivalent to $\exists \bar{X} \exists Z \forall \bar{y} \forall z \forall \bar{u} \exists w \exists v (Z \bar{y} w \wedge (Z \bar{y} z \rightarrow \varphi))$.

Thus the formula ψ used in the proof of the preceding theorem can be assumed to be Π_2 . Hence, the formula in (1) is equivalent to a Π_2 -formula and that in (2) to a Σ_2 -formula. Therefore, we have

Corollary 10.1.4 (a) MAX PB = MAX FO = MAX Π_2 .

(b) MIN PB = MIN FO = MIN Σ_2 . □

For minimization problems we can do better: We shall see in 10.1.6 that MIN PB = MIN Π_1 . In the proof, given a problem in MIN Σ_2 , we shall replace the leading existential quantifiers by relations similarly as above. The same idea underlies the following proof.

Proposition 10.1.5 MAX $\Sigma_2 = \text{MAX } \Pi_1$ and hence, MAX $\Sigma_1 \subseteq \text{MAX } \Pi_1$.

Proof. Let \mathcal{Q} be in MAX Σ_2 and let $\psi(\bar{x}, \bar{y}, \bar{Y})$ be a Π_1 -formula such that for any instance \mathcal{A} we have

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_{\bar{R}} \|\text{Sat}(\bar{R})\|,$$

where

$$\text{Sat}(\bar{R}) := \{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{R})\}.$$

We show

$$(*) \quad \max_{\bar{R}} \|\text{Sat}(\bar{R})\| = \max_{\bar{R}, P} \|\text{Sat}(\bar{R}, P)\|,$$

where $\text{Sat}(\bar{R}, P)$ is

$$\{\bar{a}\bar{b} \in A \mid \mathcal{A} \models \psi(\bar{a}, \bar{b}, \bar{R}) \wedge P\bar{a}\bar{b} \wedge \forall \bar{y}\forall \bar{z}((P\bar{a}\bar{y} \wedge P\bar{a}\bar{z}) \rightarrow \bar{y} = \bar{z})\}.$$

This gives the claim, since then $\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_{\bar{R}, P} \|\text{Sat}(\bar{R}, P)\|$, and the formula in the definition of $\text{Sat}(\bar{R}, P)$ is equivalent to a Π_1 -formula.

To prove (*), note that $\bar{a}\bar{b}_1, \bar{a}\bar{b}_2 \in \text{Sat}(\bar{R}, P)$ implies $\bar{b}_1 = \bar{b}_2$ and $\bar{a} \in \text{Sat}(\bar{R})$. Thus

$$(1) \quad \|\text{Sat}(\bar{R})\| \geq \|\text{Sat}(\bar{R}, P)\|.$$

On the other hand, for $\bar{a} \in \text{Sat}(\bar{R})$ choose a witness $\bar{b}(\bar{a})$ such that $\mathcal{A} \models \psi(\bar{a}, \bar{b}(\bar{a}), \bar{R})$ and set $P_0 := \{\bar{a}\bar{b}(\bar{a}) \mid \bar{a} \in \text{Sat}(\bar{R})\}$. Then $\text{Sat}(\bar{R}, P_0) = P_0$ and hence,

$$(2) \quad \|\text{Sat}(\bar{R})\| = \|\text{Sat}(\bar{R}, P_0)\|.$$

From (1) and (2) we conclude (*). \square

Proposition 10.1.6 $\text{MIN } \Pi_1 = \text{MIN } \Sigma_2 (= \text{MIN PB})$ and $\text{MIN } \Sigma_0 = \text{MIN } \Sigma_1$.

Proof. For the first assertion, let \mathcal{Q} be a problem in $\text{MIN } \Sigma_2$ and let $\psi(\bar{x}, \bar{y}, \bar{z}, \bar{Y})$ be a quantifier-free formula such that

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \min_{\bar{R}} \|\text{Sat}(\bar{R})\|$$

where

$$\text{Sat}(\bar{R}) := \{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y}\forall \bar{z}\psi(\bar{a}, \bar{y}, \bar{z}, \bar{R})\}.$$

We can assume that $\bar{x} = x_1 \dots x_k$ and $\bar{y} = y_1 \dots y_m$ are of the same length (if $m < k$ we extend \bar{y} by dummy variables and if $k < m$ we replace ψ by $\psi \wedge x_k = x_{k+1} = \dots = x_m$). We show

$$(*) \quad \min_{\bar{R}} \|\text{Sat}(\bar{R})\| = \min_{\bar{R}, P} \|\text{Sat}(\bar{R}, P)\|,$$

where

$$\text{Sat}(\bar{R}, P) := \{\bar{a}\bar{b} \in A \mid \mathcal{A} \models (\forall \bar{z}\psi(\bar{a}, \bar{b}, \bar{z}, \bar{R}) \wedge \neg P\bar{a}) \vee (\bar{a} = \bar{b} \wedge P\bar{a})\}.$$

This finishes the proof, since the formula in the definition of $\text{Sat}(\bar{R}, P)$ is equivalent to a Π_1 -formula.

To show (*), we verify for every \bar{R} that

$$(+) \quad \|\text{Sat}(\bar{R})\| = \min_P \|\text{Sat}(\bar{R}, P)\|.$$

Fix \bar{R} and let P be arbitrary. For $\bar{a} \in \text{Sat}(\bar{R})$ choose \bar{b} such that $\mathcal{A} \models \forall \bar{z}\psi(\bar{a}, \bar{b}, \bar{z}, \bar{R})$. If not $P\bar{a}$ then $\bar{a}\bar{b} \in \text{Sat}(\bar{R}, P)$ and if $P\bar{a}$ then $\bar{a}\bar{a} \in \text{Sat}(\bar{R}, P)$. Thus, $\|\text{Sat}(\bar{R})\| \leq \|\text{Sat}(\bar{R}, P)\|$. On the other hand, for $P = \text{Sat}(\bar{R})$, we have $\|\text{Sat}(\bar{R})\| = \|\text{Sat}(\bar{R}, P)\|$. Altogether, we see that (+), and thus (*), holds.

To prove $\text{MIN } \Sigma_0 = \text{MIN } \Sigma_1$, just omit the variables \bar{z} and the quantifiers $\forall \bar{z}$ in the preceding proof. □

By 10.1.4–10.1.6 we have

$$\begin{aligned} \text{MAX } \Sigma_0 \subseteq \text{MAX } \Sigma_1 \subseteq \text{MAX } \Pi_1 = \text{MAX } \Sigma_2 \subseteq \text{MAX } \Pi_2 = \text{MAX PB} \\ \text{MIN } \Sigma_0 = \text{MIN } \Sigma_1 \subseteq \text{MIN } \Pi_1 = \text{MIN } \Sigma_2 (= \text{MIN } \Pi_2) = \text{MIN PB}. \end{aligned}$$

All the inclusions are proper (compare [105]). The situation changes for structures with a built-in successor relation (compare [34]).

10.2 Approximable Optimization Problems

For many NPTIME optimization problems $\mathcal{Q} = (K, F, \text{cost}, \mu)$ the corresponding decision problem, in case $\mu = \max$ the set

$$\{(\mathcal{A}, k) \mid \text{there is a feasible solution } S \in F(\mathcal{A}) \text{ such that } \text{cost}(\mathcal{A}, S) \geq k\},$$

is NPTIME-complete. An example is given by MAXCUT. Unless PTIME = NPTIME, this rules out the existence of a polynomial time algorithm Π giving optimal solutions in the sense that for any instance \mathcal{A} ,

$$\Pi(\mathcal{A}) \in F(\mathcal{A}) \quad \text{and} \quad \text{opt}_{\mathcal{Q}}(\mathcal{A}) = \text{cost}(\mathcal{A}, \Pi(\mathcal{A})).$$

For practical purposes one is interested in approximation algorithms:

Definition 10.2.1 Let $0 \leq \varepsilon < 1$ and \mathcal{Q} be an NPTIME optimization problem. \mathcal{Q} is ε -*approximable* if there is a polynomial time algorithm Π , which for every instance \mathcal{A} returns a feasible solution $\Pi(\mathcal{A})$ such that

$$|\text{opt}_{\mathcal{Q}}(\mathcal{A}) - \text{cost}(\mathcal{A}, \Pi(\mathcal{A}))| \leq \varepsilon \cdot \text{opt}_{\mathcal{Q}}(\mathcal{A}).$$

That is, $\text{cost}(\mathcal{A}, \Pi(\mathcal{A})) \geq (1 - \varepsilon) \cdot \text{opt}_{\mathcal{Q}}(\mathcal{A})$ if $\mu = \max$, and $\text{cost}(\mathcal{A}, \Pi(\mathcal{A})) \leq (1 + \varepsilon) \cdot \text{opt}_{\mathcal{Q}}(\mathcal{A})$ if $\mu = \min$.

Denote by APX the class of NPTIME optimization problems that are ε -approximable for some ε with $0 \leq \varepsilon < 1$. □

To relate the class $\text{MAX } \Sigma_1$ to APX, we have to incorporate some effectivity requirements; they are fulfilled in all familiar examples.

Definition 10.2.2 A formula $\varphi(\bar{x}, \bar{Y})$ *effectively represents* the NP-TIME optimization problem \mathcal{Q} if, in addition to the conditions stated in Definition 10.1.2, there is a polynomial time algorithm that, given any instance \mathcal{A} and any relations \bar{R} on A , computes a feasible solution $S \in F(\mathcal{A})$ with

$$\text{cost}(\mathcal{A}, S) = \|\{\bar{a} \in A \mid \mathcal{A} \models \varphi(\bar{a}, \bar{R})\}\|. \quad \square$$

In the following let $\text{MAX}_e \Sigma_1$ be the class of maximization problems which can be effectively represented by a Σ_1 -formula.

Theorem 10.2.3 $\text{MAX}_e \Sigma_1 \subseteq \text{APX}$.

Proof. Consider a problem \mathcal{Q} in $\text{MAX}_e \Sigma_1$. Then, for a suitable quantifier-free formula $\psi(\bar{x}, \bar{y}, \bar{Y})$, the formula $\exists \bar{y} \psi(\bar{x}, \bar{y}, \bar{Y})$ effectively represents \mathcal{Q} ; in particular, for any instance \mathcal{A} we have

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_{\bar{R}} \|\{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{R})\}\|.$$

Set

$$A_0 := \{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{Y} \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{Y})\}.$$

Then

$$(*) \quad \|A_0\| \geq \text{opt}_{\mathcal{Q}}(\mathcal{A}).$$

Assume $\bar{y} = y_1 \dots y_n$. Since $A_0 = \{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y} \exists \bar{Y} \psi(\bar{a}, \bar{y}, \bar{Y})\}$, there is a function $f : A_0 \rightarrow A^n$ such that

$$\mathcal{A} \models \exists \bar{Y} \psi(\bar{a}, f(\bar{a}), \bar{Y})$$

holds for all $\bar{a} \in A_0$. Consider the probability space Ω of all tuples of relations \bar{R} on A (of arities corresponding to \bar{Y}) with the uniform probability distribution and let $\chi_{\mathcal{A}}$ be the random variable with

$$\chi_{\mathcal{A}}(\bar{R}) := \|\{\bar{a} \in A_0 \mid \mathcal{A} \models \psi(\bar{a}, f(\bar{a}), \bar{R})\}\|.$$

Let k be the number of atomic formulas in ψ with a relation variable in \bar{Y} . Then for the mean value $E(\chi_{\mathcal{A}})$ of $\chi_{\mathcal{A}}$ we show:

- (1) $E(\chi_{\mathcal{A}}) \geq \frac{1}{2^k} \|A_0\|$.
- (2) There is a polynomial time algorithm that, given an instance \mathcal{A} , yields relations \bar{R} on A such that

$$\chi_{\mathcal{A}}(\bar{R}) \geq E(\chi_{\mathcal{A}}).$$

This proves the claim: Given \mathcal{A} , in polynomial time we first pass to relations \bar{R} as in (2) and then, by effective representability, to a feasible solution, say $\Pi(\mathcal{A})$, with

$$\begin{aligned} \text{cost}(\mathcal{A}, \Pi(\mathcal{A})) &= \|\{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{R})\}\| \\ &\geq \|\{\bar{a} \in A_0 \mid \mathcal{A} \models \psi(\bar{a}, f(\bar{a}), \bar{R})\}\| \\ &= \chi_{\mathcal{A}}(\bar{R}) \geq \frac{1}{2^k} \|A_0\| \quad (\text{by (1) and (2)}) \\ &\geq \frac{1}{2^k} \text{opt}_{\mathcal{Q}}(\mathcal{A}) \quad (\text{by (*)}). \end{aligned}$$

To show (1), we define the indicator function I by

$$I(\bar{R}, \bar{a}) := \begin{cases} 1 & \text{if } \mathcal{A} \models \psi(\bar{a}, f(\bar{a}), \bar{R}) \\ 0 & \text{otherwise.} \end{cases}$$

Then $\chi_{\mathcal{A}}(\bar{R}) = \sum_{\bar{a} \in A_0} I(\bar{R}, \bar{a})$ and hence,

$$\begin{aligned} E(\chi_{\mathcal{A}}) &= \sum_{\bar{R} \in \Omega} \frac{1}{\|\Omega\|} \cdot \chi_{\mathcal{A}}(\bar{R}) = \frac{1}{\|\Omega\|} \sum_{\bar{R} \in \Omega} \sum_{\bar{a} \in A_0} I(\bar{R}, \bar{a}) \\ &= \sum_{\bar{a} \in A_0} \frac{1}{\|\Omega\|} \sum_{\bar{R} \in \Omega} I(\bar{R}, \bar{a}). \end{aligned}$$

Fix $\bar{a} \in A_0$. Let $\alpha_1, \dots, \alpha_k$ be the distinct atomic subformulas of ψ with a relation variable in \bar{Y} . Then \mathcal{A} determines the truth values of all atomic subformulas of ψ different from $\alpha_1, \dots, \alpha_k$. By definition of A_0 and f , there are relations \bar{P} such that $\mathcal{A} \models \psi(\bar{a}, f(\bar{a}), \bar{P})$. Thus, for the corresponding assignment of truth values to $\alpha_1, \dots, \alpha_k$, the formula ψ gets the truth value TRUE. Hence, for all relations \bar{R} leading to the same truth values of $\alpha_1, \dots, \alpha_k$, and these are at least $\frac{1}{2^k} \|\Omega\|$ many, we have $\mathcal{A} \models \psi(\bar{a}, f(\bar{a}), \bar{R})$. Thus, $\sum_{\bar{R} \in \Omega} I(\bar{R}, \bar{a}) \geq \frac{1}{2^k} \|\Omega\|$. So the equalities above give $E(\chi_{\mathcal{A}}) \geq \frac{1}{2^k} \|A_0\|$.

The last considerations show (and the same applies to similar mean values below) that $E(\chi_{\mathcal{A}})$ can be evaluated in time polynomial in $\|A\|$: First, we list A_0 and a function f of the kind in question by testing for all $\bar{a} \in A$ whether for some \bar{b} we have $\psi(\bar{a}, \bar{b}, \bar{R})$ for some \bar{R} , putting \bar{a} in A_0 , and setting $f(\bar{a}) = \bar{b}$ if \bar{b} is the first tuple with this property (note that we need \bar{R} only on the sets $\{\bar{a}, \bar{b}\}$ of limited cardinality). Now, to calculate $E(\chi_{\mathcal{A}})$, for $\sum_{\bar{R} \in \Omega} I(\bar{R}, \bar{a})$ we again need consider the behaviour of the relations \bar{R} only on the set $\{\bar{a}, f(\bar{a})\}$, the total number of the relations with the same behaviour being easily calculable.

We come to a proof of part (2). Suppose $A = \{1, \dots, n\}$. In polynomial time we construct relations \bar{R} with $\chi_{\mathcal{A}}(\bar{R}) \geq E(\chi_{\mathcal{A}})$. For this purpose we shall stepwise fix the truth value of $R\bar{e}$ for $R \in \bar{R}$ and $\bar{e} \in A$.

Let β_1, \dots, β_m be an enumeration (in a standard way) without repetitions of the elements of the set

$$\{\alpha_j(\bar{a}, f(\bar{a})) \mid 1 \leq j \leq k, \bar{a} \in A_0\}$$

(if, for example, $f(\bar{a}) = \bar{b}$, $f(\bar{b}) = \bar{a}$, $\alpha_i = R\bar{x}\bar{y}$, $\alpha_j = R\bar{y}\bar{x}$, and $i \neq j$, then $\alpha_i(\bar{a}, f(\bar{a})) = \alpha_j(\bar{b}, f(\bar{b}))$, and this instance occurs only once in β_1, \dots, β_m).

By induction on s we fix the truth value $t_s \in \{\text{TRUE}, \text{FALSE}\}$ of β_s such that the following inequality holds for the conditional expectations:

$$(*) \quad E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_s^A = t_s) \geq E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}).$$

Since $E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}) =$

$$\begin{aligned} & \frac{1}{2} E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}, \beta_s^A = \text{TRUE}) \\ & + \frac{1}{2} E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}, \beta_s^A = \text{FALSE}), \end{aligned}$$

in order to fix t_s we only have to check whether

$$\begin{aligned} & E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}, \beta_s^A = \text{TRUE}) \\ & \geq E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{s-1}^A = t_{s-1}, \beta_s^A = \text{FALSE}), \end{aligned}$$

and by the remarks above this can be done in polynomial time. For $s = m$ we have by (*)

$$\begin{aligned} E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_m^A = t_m) & \geq E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_{m-1}^A = t_{m-1}) \\ & \geq \dots \geq E(\chi_{\mathcal{A}} | \beta_1^A = t_1) \geq E(\chi_{\mathcal{A}}). \end{aligned}$$

Since for relations \bar{R} and \bar{R}' that coincide on all atomic subformulas of $\psi(\bar{a}, f(\bar{a}), \bar{Y})$ with $\bar{a} \in A_0$, we have

$$\chi_{\mathcal{A}}(\bar{R}) = \chi_{\mathcal{A}}(\bar{R}'),$$

we see that for any \bar{R} with truth values t_1, \dots, t_m on β_1, \dots, β_m , respectively, we have

$$\chi_{\mathcal{A}}(\bar{R}) = E(\chi_{\mathcal{A}} | \beta_1^A = t_1, \dots, \beta_m^A = t_m) \quad (\geq E(\chi_{\mathcal{A}})).$$

Thus we obtain the relations \bar{R} as claimed in (2) by fixing the values on β_1, \dots, β_m according to t_1, \dots, t_m and arbitrarily (say, as FALSE) on all other tuples. \square

In the first part of the preceding proof we have shown

Corollary 10.2.4 *Assume \mathcal{Q} is in $\text{MAX}_e \Sigma_1$, say, for every instance \mathcal{A} we have*

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_{\bar{R}} \|\{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{R})\}\|$$

with quantifier-free $\psi(\bar{x}, \bar{y}, \bar{Y})$. Then there is a polynomial time algorithm that, given an instance \mathcal{A} , yields relations $\Pi(\mathcal{A})$ on A such that

$$\text{cost}(\mathcal{A}, \Pi(\mathcal{A})) \geq \frac{1}{2^k} \|A_0\|,$$

where k is the number of atomic formulas in ψ with a relation variable in \bar{Y} and where $A_0 := \{\bar{a} \in A \mid \mathcal{A} \models \exists \bar{Y} \exists \bar{y} \psi(\bar{a}, \bar{y}, \bar{Y})\}$. \square

As an application we obtain

Proposition 10.2.5 MAXCUT is $\frac{1}{2}$ -approximable.

Proof. If \mathcal{Q} denotes MAXCUT, we have for any graph \mathcal{A} ,

$$(+)\quad \text{opt}_{\mathcal{Q}}(\mathcal{A}) = \max_R \|\{(a, b) \mid \mathcal{A} \models Eab \wedge Ra \wedge \neg Rb\}\|.$$

Thus, the preceding corollary with $\psi = Exy \wedge Yx \wedge \neg Yy$, \bar{y} the empty sequence, $k = 2$, and $A_0 = \{(a, b) \mid \mathcal{A} \models \exists Y(Eab \wedge Ya \wedge \neg Yb)\} = E^A$ yields a polynomial time algorithm Π such that

$$\text{cost}(\mathcal{A}, \Pi(\mathcal{A})) \geq \frac{1}{4} \|E^A\|.$$

Since $E^A ab$ implies $E^A ba$, we have by (+),

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) \leq \frac{1}{2} \|E^A\|.$$

Altogether,

$$\text{cost}(\mathcal{A}, \Pi(\mathcal{A})) \geq \frac{1}{2} \text{opt}_{\mathcal{Q}}(\mathcal{A}). \quad \square$$

The following two observations show that possible extensions of Theorem 10.2.3 fail.

Denote by PTAS the class of NPTIME optimization problems \mathcal{Q} with a polynomial time approximation scheme, i.e., with the property that there is an effective procedure Π that returns, for every rational ϵ with $0 < \epsilon < 1$, a polynomial time algorithm Π_ϵ witnessing that \mathcal{Q} is ϵ -approximable. It has been shown that, unless PTIME = NPTIME, MAXCUT \notin PTAS; in particular, as MAXCUT \in MAX $_\epsilon$ Σ_1 , we have MAX $_\epsilon$ $\Sigma_1 \not\subseteq$ PTAS.

We come to the second observation. Let \mathcal{Q} be the problem MAX-CLIQUE: Given a graph \mathcal{G} it asks for the maximal size of a clique in \mathcal{G} . Hence, MAXCLIQUE \in MAX $_\epsilon$ Π_1 , since

$$\text{opt}_{\mathcal{Q}}(\mathcal{G}) = \max_R \|\{a \in G \mid \mathcal{G} \models Ra \wedge \forall x \forall y ((Rx \wedge Ry \wedge x \neq y) \rightarrow Exy)\}\|.$$

It is known that, unless PTIME = NPTIME, MAXCLIQUE \notin APX, showing that Theorem 10.2.3 cannot be extended to MAX $_\epsilon$ Π_1 .

We close this section with two “positive” remarks on Theorem 10.2.3. The first one gives a generalization, the second one shows how to enlarge its scope of applicability.

We can change the definition of a first-order definable optimization problem by considering in 10.1.2 formulas $\varphi(\bar{x}, \bar{u}, \bar{Y})$ and requiring that

$$\text{opt}_{\mathcal{Q}}(\mathcal{A}) = \frac{\mu}{\bar{b}, \bar{R}} \|\{\bar{a} \in A \mid \mathcal{A} \models \varphi[\bar{a}, \bar{b}, \bar{R}]\}\|$$

(“optimization over elements and relations”). Replacing elements by relations, we see that this does not change the notion of a first-order definable optimization problem; however, the classes $\text{MAX } \Sigma_1$ and $\text{MAX}_e \Sigma_1$ strictly increase, if we allow optimization over elements and relations (note that the replacement of constants by relation symbols uses universal quantifiers). Nevertheless, Theorem 10.2.3 remains true for the larger class $\text{MAX}_e \Sigma_1$ (indeed with essentially the same proof).

The following example shows how one can enlarge the scope of applicability of the preceding theorem. Consider the optimization problem known as **MAXIMUM CONNECTED COMPONENT (MCC)**. Given a graph, it asks for the maximum of the cardinalities of connected components. Since the connected components are computable in polynomial time, there is a polynomial time algorithm solving the optimization problem. By 10.1.4(a) we know that $\text{MCC} \in \text{MAX } \Pi_2$. On the other hand, $\text{MCC} \notin \text{MAX } \Pi_1$ (see Exercise 10.2.6 below) and hence, $\text{MCC} \notin \text{MAX } \Sigma_1$ by 10.1.5. Therefore, our logical approach does not yield $\text{MCC} \in \text{APX}$, even though there is a polynomial time algorithm giving an optimal solution.

However, since the transitive closure is computable in polynomial time, the approximability of MCC is equivalent to the approximability of MCC' , where MCC' is the problem

Given $\mathcal{G} = (G, E^G, R^G)$ with $R^G = \text{TC}(E^G)$, a “graph with transitive closure”, compute the maximal cardinality of its connected components.

Clearly, MCC' is in $\text{MAX}_e \Sigma_1$ (with maximization over elements) since

$$\text{opt}_{\text{MCC}'}(\mathcal{A}) = \max_b \|\{a \mid \mathcal{G} \models Rab\}\|.$$

Now, 10.2.3 (with maximization over elements) yields $\text{MCC}' \in \text{APX}$.

Of course, the same idea can be applied to any other optimization problem: The values of polynomial time global relations can be added to the instances by free without changing the behaviour of the optimization problem with respect to the notion of approximability we have considered. In [11] this extra relations have been “added” on the logical side, allowing formulas of FO(LFP) instead of FO .

Exercise 10.2.6 (a) Give a direct proof that $\text{MCC} \in \text{MAX } \Pi_2$. (b) Show that $\text{MCC} \notin \text{MAX } \Pi_1$ (even not with maximization over elements). \square

Notes 10.2.7 The main references for this chapter are [104, 105, 129]; they also contain corresponding historical remarks.

11. Logics for PTIME

There are logics (for instance FO(IFP); cf. Chapter 7) that capture PTIME on ordered structures. Instantly, this result evokes the question whether it can be strengthened to the case of not necessarily ordered structures:

Is there a logic strongly capturing PTIME?

There are examples of such logics (cf. 11.1.2); however, they are quite artificial, at the same time showing that we should incorporate some obvious requirements of effectivity in our notion of strongly capturing. We thus are led to an effective notion of a logic strongly capturing PTIME and, therefore, may reformulate the question above as follows:

(*) *Is there a logic effectively strongly capturing PTIME?*

The question is considered to be the most prominent open problem in finite model theory and, at the same time, a difficult one. For instance, a proof that such a logic does not exist, would yield that $\text{PTIME} \neq \text{NPTIME}$ (cf. section 11.1) and hence, solve the most prominent open problem in complexity theory. Moreover, in section 11.1 we show that the question (*) is strongly related to the isomorphism problem for finite structures.

The statement “FO(IFP) captures PTIME” can be reformulated as “FO(IFP) effectively strongly captures PTIME on the class of ordered structures”. In section 11.2 we are led to investigate specific classes of (not necessarily ordered) structures whether they allow for a similar result, i.e., we are led to search for partial solutions of (*). We present some positive results.

The discussion of specific logics that may be candidates for a positive solution of (*) itself, is transferred to the next chapter. There we shall also see that if there is any logic at all that effectively strongly captures PTIME, then there is such a logic with a familiar syntax.

In the following not only structures, but also formulas and Turing machines will be considered as objects of computations. When doing so, we shall tacitly assume that vocabularies, formulas, and machines are coded by $\{0, 1\}$ -words in some reasonable way.

11.1 Logics and Invariants

First, we recall the definition of a logic \mathcal{L} capturing a complexity class \mathcal{C} and of \mathcal{L} strongly capturing \mathcal{C} .

For a class K of structures we introduced the class $K_{<}$ of its ordered representations,

$$K_{<} = \{(\mathcal{A}, <) \mid \mathcal{A} \in K, < \text{ an ordering of } A\}.$$

In particular, if $<_1$ and $<_2$ are orderings of A then

$$(*) \quad (\mathcal{A}, <_1) \in K_{<} \quad \text{iff} \quad (\mathcal{A}, <_2) \in K_{<}.$$

The logic \mathcal{L} captures the complexity class \mathcal{C} iff for every class K of *ordered* structures,

$$K \in \mathcal{C} \quad \text{iff} \quad K \text{ is axiomatizable in } \mathcal{L}.$$

And \mathcal{L} strongly captures \mathcal{C} iff for every class K of (not necessarily ordered) structures,

$$K_{<} \in \mathcal{C} \quad \text{iff} \quad K \text{ is axiomatizable in } \mathcal{L}.^1$$

Already several times we mentioned the following open question: Is there a logic strongly capturing PTIME? Below we present two examples of such logics which will turn out to be “unsatisfactory”, thus leading to a sharper version of the question.

Recall that a sentence φ of vocabulary $\tau_{<} = \tau \dot{\cup} \{<\}$ (i.e., $< \notin \tau$) is *order-invariant in the finite* if

$$(\mathcal{A}, <_1) \models \varphi \quad \text{iff} \quad (\mathcal{A}, <_2) \models \varphi$$

holds for all (finite) τ -structures \mathcal{A} and orderings $<_1, <_2$ of A .

By (*) above, if $K_{<} = \text{Mod}(\varphi)$ then φ is order-invariant in the finite.

Exercise 11.1.1 For sufficiently rich τ the set

$$\{\varphi \in \text{FO}[\tau_{<}] \mid \varphi \text{ a sentence order-invariant in the finite}\}$$

is not decidable and hence, not recursively enumerable. Hint: Let σ be a vocabulary such that

$$\{\psi \in \text{FO}[\sigma] \mid \psi \text{ a sentence unsatisfiable in the finite}\}$$

¹ Note that on the left side of the last two equivalences, K and $K_{<}$ are classes of ordered structures. This is an accordance with the conventions we agreed upon in Chapter 7, that only classes of ordered structures are considered as members of complexity classes.

is not decidable (cf. Trahtenbrot’s Theorem 7.2.1). Let $\tau := \sigma \cup \{P\}$, where P is a unary relation symbol, $P \notin \sigma$. For a sentence $\psi \in \text{FO}[\sigma]$ let φ_ψ be the $\text{FO}[\tau_{<}]$ -sentence

$$\varphi_\psi := (\psi \rightarrow \exists x(\text{“}x \text{ is the } <\text{-minimal element”} \wedge Px)).$$

Then, φ_ψ is order-invariant in the finite iff $\psi \wedge \exists x \exists y x \neq y$ is unsatisfiable in the finite. \square

We present the promised examples.

Example 11.1.2 Let \mathcal{L}_1 be the logic whose sentences of vocabulary τ are the $\text{FO}(\text{IFP})[\tau_{<}]$ -sentences, where the satisfaction relation $\mathcal{A} \models^{\mathcal{L}_1} \varphi$ for a τ -structure \mathcal{A} and an $\mathcal{L}_1[\tau]$ -sentence φ is defined by

$$\mathcal{A} \models^{\mathcal{L}_1} \varphi \text{ iff } \varphi \text{ is order-invariant in the finite and there is an ordering } <^A \text{ on } A \text{ such that } (\mathcal{A}, <^A) \models^{\text{IFP}} \varphi$$

(on the right side, $(\mathcal{A}, <^A) \models^{\text{IFP}} \varphi$ means that $(\mathcal{A}, <^A)$ is a model of φ in the semantics of $\text{FO}(\text{IFP})$).

Let \mathcal{L}_2 be the logic whose sentences of vocabulary τ are the $\text{FO}(\text{IFP})[\tau_{<}]$ -sentences order-invariant in the finite where now the satisfaction relation is defined by

$$\mathcal{A} \models^{\mathcal{L}_2} \varphi \text{ iff there is an ordering } <^A \text{ on } A \text{ such that } (\mathcal{A}, <^A) \models^{\text{IFP}} \varphi.$$

As $\text{FO}(\text{IFP})$ captures PTIME, one easily verifies that both \mathcal{L}_1 and \mathcal{L}_2 strongly capture PTIME. However, both logics have unsatisfactory features: Concerning \mathcal{L}_1 , in view of 11.1.1 there is no effective procedure assigning to every \mathcal{L}_1 -sentence φ an algorithm that, given an ordered version of \mathcal{A} , evaluates $\mathcal{A} \models \varphi$ in time polynomial in $\|\mathcal{A}\|$. On the other hand, given an \mathcal{L}_2 -sentence φ , there is such an effective procedure. But now, using 11.1.1 again, we see that for sufficiently rich τ the set of $\mathcal{L}_2[\tau]$ -sentences is not decidable. \square

None of the logics just considered satisfies at the same time that

- the set of sentences is decidable
- there is an effective procedure assigning to every φ a polynomially time-bounded algorithm testing $\mathcal{A} \models \varphi$.

We incorporate both requirements in our notion of capturing given by the following definition.² When speaking of complexity classes, we have in mind one of the concrete complexity classes considered so far, and we think of a logic \mathcal{L} as given – for any vocabulary τ – by a set $\mathcal{L}[\tau]_0$, the set of \mathcal{L} -sentences of vocabulary τ , and by a relation $\models^{\mathcal{L},\tau}$ between τ -structures and \mathcal{L} -sentences of vocabulary τ , the \mathcal{L} -satisfaction relation for τ .

² The reader may wonder why, say, in the second requirement we do not ask for a polynomial time procedure instead of an effective procedure. In fact, there is no real difference: By standard techniques one can pass from a logic satisfying the weaker conditions to a logic satisfying the stronger ones.

Definition 11.1.3 Let \mathcal{L} be a logic and \mathcal{C} a complexity class. \mathcal{L} *effectively strongly captures* \mathcal{C} , $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$, if

- \mathcal{L} strongly captures \mathcal{C} ;
- for every vocabulary τ
 - (i) $\mathcal{L}[\tau]_0$ is a decidable set;
 - (ii) there is an effective procedure that to every sentence $\varphi \in \mathcal{L}[\tau]_0$ assigns a pair (M, f) , where M is a Turing machine that accepts $\text{Mod}(\varphi)_{<}$ and f is (the code of) a function witnessing that M is resource-bounded according to \mathcal{C} .³ □

In all concrete cases considered in Chapter 7 we have implicitly proved that the logics and the corresponding complexity classes satisfy the effectivity conditions (i) and (ii). For example, the proof of 7.4.2(a) yields a procedure that to every $\varphi \in \text{FO}(\text{IFP})[\tau_{<}]$ assigns a pair (M, d) where M is an x^d time-bounded deterministic Turing machine accepting the class of ordered models of φ .

In the definition above we do not impose limitations on the syntax of a logic (besides (i)). For example, NPTIME is not only effectively strongly captured by the logic of Σ_1^1 -sentences, but also by the logic \mathcal{L} with

- $\mathcal{L}[\tau]_0 := \{(M, d) \mid d \geq 1 \text{ and } M \text{ is a nondeterministic Turing machine for ordered representations of } \tau\text{-structures}\}$
- for every $\varphi = (M, d) \in \mathcal{L}[\tau]_0$ and every τ -structure \mathcal{A} ,

$\mathcal{A} \models \varphi$ iff M accepts some ordered representation of \mathcal{A} in $\leq \|\mathcal{A}\|^d$ steps.

This logic is very artificial and lacks the logical flavour we are accustomed to. Moreover, it is tied up so closely with the definition of NPTIME that the characterization of NPTIME in terms of this logic does not give us any additional insight. We shall show in 12.3.17 that in case PTIME is effectively strongly captured by any logic at all, it can already be captured by a logic with a familiar syntax.

Altogether, we reformulate our main question as

- (*) *Is there a logic effectively strongly capturing PTIME?*

In the remainder of this section we study the relationship of (*) with other problems, thereby presenting methods and results that may be of help in getting an answer.

Proposition 11.1.4 *If the answer to (*) is “no” then $\text{PTIME} \neq \text{NPTIME}$.*

³ To be definite, in case $\mathcal{C} \in \{\text{LOGSPACE}, \text{NLOGSPACE}, \text{PSPACE}\}$ we code f by a natural number d meaning that M is d -log space-bounded or (for PSPACE) x^d space-bounded; in case $\mathcal{C} \in \{\text{PTIME}, \text{NPTIME}\}$ we also code f by a number d meaning that M is x^d time-bounded.

Proof. There seems to be a very short proof: Assume $\text{PTIME} = \text{NPTIME}$. Since $\Sigma_1^1 \equiv_{\text{es}} \text{NPTIME}$, we have $\Sigma_1^1 \equiv_{\text{es}} \text{PTIME}$. However, this argument has a gap: By $\Sigma_1^1 \equiv_{\text{es}} \text{NPTIME}$ we know that, given τ , there is an effective procedure assigning to every $\Sigma_1^1[\tau]$ -sentence φ a number $d \geq 1$ and a nondeterministic x^d time-bounded Turing machine M accepting $\text{Mod}(\varphi)_{<}$. But how do we get a *deterministic* machine effectively?

To close this gap one shows, using an NPTIME -complete problem, that $\text{PTIME} = \text{NPTIME}$ implies an effective version of $\text{PTIME} = \text{NPTIME}$. More precisely: Let SAT be the set of satisfiable propositional formulas in conjunctive normal form. SAT is NPTIME -complete (cf. [87]). Even more: There is a polynomial time algorithm A_0 that assigns to every $d \geq 1$, every nondeterministic x^d time-bounded Turing machine M , and every input string u for M a propositional formula $\alpha(u, (M, d))$ in conjunctive normal form such that

$$(+) \quad M \text{ accepts } u \quad \text{iff} \quad \alpha(u, (M, d)) \in \text{SAT}.$$

Assume $\text{PTIME} = \text{NPTIME}$ and let A_1 be a deterministic polynomially time-bounded algorithm for SAT . By (+), using the algorithms A_0 and A_1 and time bounds for them, one can effectively assign to every $d \geq 1$ and nondeterministic x^d time-bounded M a number $d' \geq 1$ and a deterministic $x^{d'}$ time-bounded algorithm accepting the same language as M . \square

When defining the logics in Example 11.1.2, we essentially restricted ourselves to sentences order-invariant in the finite to ensure that the model classes $\text{Mod}^{\mathcal{L}_1}(\varphi)$ and $\text{Mod}^{\mathcal{L}_2}(\varphi)$ belong to PTIME . We could have dispensed with this restriction if there would exist a PTIME -algorithm defining a “canonical” ordered version of every structure. Proposition 11.1.6 contains the precise statement and the following Definition 11.1.5 gives the precise notion.

Denote by Str and $\text{Str}[\tau]$ the class of all finite structures and of all finite τ -structures, respectively.

Definition 11.1.5 A *PTIME-canonization* \mathbf{C} consists of functions

$$C_\tau : \text{Str}[\tau] \rightarrow \text{Str}[\tau_{<}]$$

for every vocabulary τ such that:

- (1) For all $\mathcal{A} \in \text{Str}[\tau]$, $(\mathcal{A} \cong C_\tau(\mathcal{A})|_\tau$ and $<^{C_\tau(\mathcal{A})}$ is an ordering).⁴
- (2) For all $\mathcal{A}, \mathcal{B} \in \text{Str}[\tau]$, $\mathcal{A} \cong \mathcal{B}$ implies $C_\tau(\mathcal{A}) \cong C_\tau(\mathcal{B})$.
- (3) C_τ is PTIME -computable, more precisely: there is a PTIME -algorithm that, applied to $(\mathcal{A}, <^{\mathcal{A}}) \in \text{Str}[\tau]_{<}$, gives (the encoding of) the ordered structure $C_\tau(\mathcal{A})$.⁵ \square

⁴ Recall that $C_\tau(\mathcal{A})|_\tau$ denotes the τ -reduct of $C_\tau(\mathcal{A})$.
⁵ In Chapter 7 we agreed upon how to regard ordered structures as inputs to a Turing machine M . By those conventions it is clear that for isomorphic $(\mathcal{A}, <^{\mathcal{A}})$

Clearly, in view of (1), condition (2) is equivalent to
 (2') For all $\mathcal{A}, \mathcal{B} \in \text{Str}[\tau]$, $\mathcal{A} \cong \mathcal{B}$ iff $C_\tau(\mathcal{A}) \cong C_\tau(\mathcal{B})$.

Proposition 11.1.6 *If there exists a PTIME-canonization then there is a logic effectively strongly capturing PTIME.*

Proof. Let \mathbf{C} be a PTIME-canonization. Consider the logic \mathcal{L} whose sentences of vocabulary τ are the FO(IFP)[$\tau_{<}$]-sentences. The satisfaction relation is defined by

$$(*) \quad \mathcal{A} \models^{\mathcal{L}, \tau} \varphi \quad \text{iff} \quad C_\tau(\mathcal{A}) \models^{\text{IFP}} \varphi.$$

\mathcal{L} effectively strongly captures PTIME: Clearly, the set of $\mathcal{L}[\tau]$ -sentences is decidable. From Chapter 7 we know that there is an effective procedure assigning to every FO(IFP)[$\tau_{<}$]-sentence φ a pair (M_0, d_0) , where M_0 is an x^{d_0} time-bounded Turing machine accepting the class of ordered models of φ . Therefore, using a PTIME algorithm for C_τ , the equivalence (*) shows that one effectively can assign to every $\mathcal{L}[\tau]$ -sentence φ a pair (M, d) , where M is an x^d time-bounded Turing machine accepting $\text{Mod}^{\mathcal{L}}(\varphi)_{<}$. In particular, $\text{Mod}^{\mathcal{L}}(\varphi)_{<}$ is in PTIME.

Conversely, if K is a class of τ -structures with $K_{<} \in \text{PTIME}$, then $K_{<} = \text{Mod}^{\text{FO(IFP)}}(\varphi)$ for some FO(IFP)[$\tau_{<}$]-sentence φ . But then, $K = \text{Mod}^{\mathcal{L}}(\varphi)$. □

For an ordered $\tau_{<}$ -structure \mathcal{A} , let \mathcal{A}^+ be its numerical representation, i.e., \mathcal{A}^+ is the unique $\tau_{<}$ -structure isomorphic to \mathcal{A} with $A^+ = \{0, \dots, \|\mathcal{A}\| - 1\}$ and with $<^{\mathcal{A}^+}$ the natural ordering on A^+ . Assume that \mathbf{C} is a PTIME-canonization. By passing to $(C_\tau(\mathcal{A}))^+$, one sees that (2) (compare (2'), too) in the definition of PTIME-canonization can be replaced by

$$\text{For all } \mathcal{A}, \mathcal{B} \in \text{Str}[\tau], \quad \mathcal{A} \cong \mathcal{B} \quad \text{iff} \quad C_\tau(\mathcal{A}) = C_\tau(\mathcal{B}).$$

The structures $(C_\tau(\mathcal{A}))^+$ can be coded in a canonical way, say, by words over the alphabet $\{0, 1\}$. Thus, a canonization leads to an “invariantization”, i.e., to a map associating with every structure \mathcal{A} a word only depending on the isomorphism type of \mathcal{A} . Again, we give a precise definition and statement.

Definition 11.1.7 A PTIME-invariantization \mathbf{I} consists of functions

$$I_\tau : \text{Str}(\tau) \rightarrow \{0, 1\}^*$$

such that for every vocabulary τ we have:

and $(\mathcal{B}, <^{\mathcal{B}})$ the result of M applied to $(\mathcal{A}, <^{\mathcal{A}})$ is the same as that of M applied to $(\mathcal{B}, <^{\mathcal{B}})$. Therefore, if C_τ is computable then

$$\mathcal{A} \cong \mathcal{B} \quad \text{implies} \quad C_\tau(\mathcal{A}) \cong C_\tau(\mathcal{B}).$$

Thus, in the above definition, property (2) is redundant. Nevertheless, here and in the following, we stay with such redundancies in order to emphasize both the structural and the computational aspects of such functions.

- (1) For all $\mathcal{A}, \mathcal{B} \in \text{Str}[\tau]$, $\mathcal{A} \cong \mathcal{B}$ iff $I_\tau(\mathcal{A}) = I_\tau(\mathcal{B})$.
- (2) I_τ is PTIME-computable (more precisely: there is a PTIME-algorithm that, applied to $(\mathcal{A}, <^A) \in \text{Str}[\tau]_{<}$, gives the word $I_\tau(\mathcal{A})$).

Proposition 11.1.8 *If there is a PTIME-canonization, then there is a PTIME-invariantization.*

Proof. If \mathbf{C} is a PTIME-canonization, then let $I_\tau(\mathcal{A})$ be the $\{0, 1\}$ -code of $(C_\tau(\mathcal{A}))^+$. By the remarks above, \mathbf{I} is a PTIME-invariantization. \square

It may be surprising that the converse of the preceding proposition is also true, that is, an “arbitrary” invariant already leads to a canonical ordering.

Theorem 11.1.9 *If there is a PTIME-invariantization, then there is a PTIME-canonization.*

Proof. Assume that \mathbf{I} is a PTIME-invariantization. Let τ be an arbitrary vocabulary and let $<$ be a new binary relation symbol. Set $\sigma = \tau_{<}$. First, using $I_\sigma : \text{Str}(\sigma) \rightarrow \{0, 1\}^*$, for every τ -structure \mathcal{A} and every ordering \prec^A on A , we stepwise define an ordering $<^A$ on A .

Fix \mathcal{A} (with at least two elements) and an ordering \prec^A on A . If a_1, \dots, a_l are distinct elements of A , let

$$[a_1, \dots, a_l] := \{(a_i, a_j) \mid 1 \leq i < j \leq l\}$$

if $l \geq 2$ and $[a_1] := \{(a_1, a_1)\}$. We start the step by step construction of $<^A$: Let w_1 be the first element in the lexicographic ordering of $\{0, 1\}^*$ of the set

$$\{I_\sigma((\mathcal{A}, [a])) \mid a \in A\}.$$

Among the a 's with $w_1 = I_\sigma((\mathcal{A}, [a]))$, let a_1 be the \prec^A -first one.

In the second step let w_2 be the first element in

$$\{I_\sigma((\mathcal{A}, [a_1, a])) \mid a \in A, a \neq a_1\}.$$

Among the a 's with $w_2 = I_\sigma((\mathcal{A}, [a_1, a]))$ and $a \neq a_1$, let a_2 be the \prec^A -first one.

After n steps, where $n := \|A\|$, we have obtained a_1, \dots, a_n with $A = \{a_1, \dots, a_n\}$. We set

$$<^A := [a_1, \dots, a_n]$$

and define \mathbf{C} by setting

$$C_\tau(\mathcal{A}) := \text{the numerical representation of } (\mathcal{A}, <^A).$$

To see that \mathbf{C} is a PTIME-canonization, we first observe that the transition from (\mathcal{A}, \prec^A) to $(\mathcal{A}, <^A)$ can effectively be performed in polynomial time.

Trivially, $\mathcal{A} \cong (\mathcal{A}, <^A)|_\tau$ and hence, $\mathcal{A} \cong C_\tau(\mathcal{A})|_\tau$. Finally, suppose that $\mathcal{B} \cong \mathcal{A}$ and \prec^B is an ordering of B . Let $<^B = [b_1, \dots, b_n]$ be the ordering given by the stepwise procedure above applied to (\mathcal{B}, \prec^B) . Then, for $l = 1, \dots, n$,

$$(+) \quad (\mathcal{A}, [a_1, \dots, a_l]) \cong (\mathcal{B}, [b_1, \dots, b_l]).$$

For $l = n$ this yields $(\mathcal{A}, <^A) \cong (\mathcal{B}, <^B)$ and hence, $C_\tau(\mathcal{A}) = C_\tau(\mathcal{B})$.

The proof of (+) proceeds by induction on l . We show the case $l = 1$: Suppose $\pi : \mathcal{A} \cong \mathcal{B}$. Then, for all $a \in A$,

$$\pi : (\mathcal{A}, [a]) \cong (\mathcal{B}, [\pi(a)]).$$

Thus, by property (1) of the invariantization **I**,

$$\{I_\sigma((\mathcal{A}, [a]) \mid a \in A) \} = \{I_\sigma((\mathcal{B}, [b]) \mid b \in B)\},$$

and hence, $I_\sigma((\mathcal{A}, [a_1])) = I_\sigma((\mathcal{B}, [b_1]))$. Again by property (1), we therefore have $(\mathcal{A}, [a_1]) \cong (\mathcal{B}, [b_1])$. \square

Together with 11.1.6 the theorem yields:

Corollary 11.1.10 *If there is a PTIME-invariantization, then there is a logic effectively strongly capturing PTIME.* \square

Of course, the notion of canonization can be defined for any deterministic complexity class \mathcal{C} : A \mathcal{C} -canonization **C** consists of functions $C_\tau : \text{Str}[\tau] \rightarrow \text{Str}[\tau_{<}]$ satisfying (1), (2) of Definition 11.1.5 and

(3) C_τ is \mathcal{C} -computable (more precisely: there is an algorithm according to \mathcal{C} that, applied to $(\mathcal{A}, <^A) \in \text{Str}[\tau]_{<}$, gives (the encoding of) the ordered structure $C_\tau(\mathcal{A})$).⁶

Exercise 11.1.11 Generalizing the proof of 11.1.6 show: If there are a \mathcal{C} -canonization and a logic effectively strongly capturing \mathcal{C} on ordered structures then there is a logic \mathcal{L} with $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$. \square

Exercise 11.1.12 Show the existence of a PSPACE-canonization. Hint: Given a τ -structure \mathcal{A} , let $C_\tau(\mathcal{A})$ be the numerical representation of $(\mathcal{A}, <)$, where $<$ is a minimal order on \mathcal{A} (compare 8.5.13). \square

We conclude with some remarks that shed still further light on the main question.

Suppose there is a logic \mathcal{L} effectively strongly capturing the complexity class \mathcal{C} . Then, given τ , we can effectively enumerate the $\mathcal{L}[\tau]$ -sentences. Using the effective transition from sentences φ to pairs (M, f) , where M is a Turing machine accepting $\text{Mod}(\varphi)_{<}$ and f a resource bound of M according to \mathcal{C} , we get that \mathcal{C} is effectively enumerable in the sense of the following

⁶ For $\mathcal{C} = \text{LOGSPACE}$ this means that C_τ is computable by a log space-bounded transducer, cf. section 7.5.1.

Definition 11.1.13 Let \mathcal{C} be a complexity class. An *effective enumeration* \mathbf{F} of \mathcal{C} consists of computable functions F_τ , τ a vocabulary, that are defined on the set \mathbb{N} of natural numbers, with the properties

- (1) For any $i \in \mathbb{N}$, $F_\tau(i)$ is a pair (M, f) where M is a Turing machine accepting a class $K_<$ with $K \subseteq \text{Str}[\tau]$, and f witnesses that M is according to \mathcal{C} .
- (2) For any class $K \subseteq \text{Str}[\tau]$ such that $K_<$ is in \mathcal{C} , there is an i with $F_\tau(i) = (M, f)$ such that the machine M accepts $K_<$. □

By the remark preceding the definition we have the direction from (i) to (ii) of

Proposition 11.1.14 *For a complexity class \mathcal{C} , the following are equivalent:*

- (i) *There is a logic effectively strongly capturing \mathcal{C} .*
- (ii) *There is an effective enumeration of \mathcal{C} .*

Proof. We have to show that (ii) implies (i). Let \mathbf{F} be an effective enumeration of \mathcal{C} . First assume that, for all τ , the set $\{F_\tau(i) \mid i \in \mathbb{N}\}$ is decidable. Define the logic \mathcal{L} as follows. The set of $\mathcal{L}[\tau]$ -sentences is $\{F_\tau(i) \mid i \in \mathbb{N}\}$. And if $F_\tau(i) = (M, f)$ where M accepts the class $K_<$, then

$$A \models^{\mathcal{L}} F_\tau(i) \quad \text{iff} \quad A \in K.$$

Hence, $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$.

Otherwise, we replace \mathbf{F} by an enumeration \mathbf{F}^* such that, for all τ , the set $\{F_\tau^*(i) \mid i \in \mathbb{N}\}$ is decidable. We define $F_\tau^*(i)$ as follows. Let $F_\tau(i) = (M, f)$. Set $F_\tau^*(i) = (M^*, f)$, where M^* behaves the same way as M , however, the length of (the coding of) M^* is greater than i ; this is achieved, say, by adding to M i -many new instructions that never will be used. To decide whether a pair (M', f') belongs to $\{F_\tau^*(i) \mid i \in \mathbb{N}\}$, we check whether (M', f') coincides with some $F_\tau^*(i)$ for $1 \leq i \leq \text{length of } M'$. □

By the proposition, the complexity classes above PTIME such as NPTIME or PSPACE are effectively enumerable. If we view an effective enumeration as providing a tool for a systematic exploration of the complexity class concerned, NPTIME and PSPACE are explorable in this sense, whereas it is open whether PTIME or its subclasses LOGSPACE and NLOGSPACE are explorable.

11.2 PTIME on Classes of Structures

We have already remarked that the question (“main question”) whether there is a logic effectively strongly capturing PTIME, has a positive answer if we restrict ourselves to ordered structures, FO(IFP) being such a logic. We may

consider this result as a partial answer to the main question and, thus, are led to a more systematic study of such restrictions. In the following we show, for example, that the main question is equivalent to the question whether there is a logic effectively strongly capturing the PTIME properties on the class GRAPH of graphs. Moreover, we shall see that for certain subclasses of GRAPH such as the class of trees the corresponding question has a positive answer.

If S is a class of structures, $S \subseteq \text{Str}$, and τ a vocabulary, let $S[\tau] := \{\mathcal{A} \in S \mid \mathcal{A} \text{ a } \tau\text{-structure}\}$.

Definition 11.2.1 Let \mathcal{L} be a logic, \mathcal{C} a complexity class, and S a class of structures, $S \subseteq \text{Str}$. \mathcal{L} *effectively strongly captures* \mathcal{C} on S , $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S , if for all τ

- (1) $\mathcal{L}[\tau]_0$ is decidable;
- (2) $\{K \cap S[\tau] \mid K \subseteq \text{Str}[\tau], K_{<} \in \mathcal{C}\} = \{\text{Mod}^{\mathcal{L}}(\varphi) \cap S[\tau] \mid \varphi \in \mathcal{L}[\tau]_0\}$;⁷
- (3) there is an effective procedure that to every $\varphi \in \mathcal{L}[\tau]_0$ assigns a pair (M, f) , where f is a function witnessing that M is resource-bounded according to \mathcal{C} and where for every $(\mathcal{A}, <) \in S[\tau]_{<}$ the machine M accepts $(\mathcal{A}, <)$ iff $\mathcal{A} \models \varphi$. □

Clearly, if S and S' are classes of structures with $S \subseteq S' \subseteq \text{Str}$ and if $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S' , then $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S .

In many concrete situations we have $S[\tau]_{<} \in \mathcal{C}$. Then, condition (2) of the previous definition can be simplified:

Proposition 11.2.2 Let \mathcal{L} be a logic, \mathcal{C} a complexity class and S a class of structures with $S[\tau]_{<} \in \mathcal{C}$ for all τ .

- (a) If \mathcal{L} is a logic that for all τ satisfies (1), (3) of Definition 11.2.1 and

(2*) for all $K \subseteq S[\tau]$,

$$K_{<} \in \mathcal{C} \quad \text{iff} \quad \text{there is } \varphi \in \mathcal{L}[\tau]_0 \text{ such that } K = \text{Mod}^{\mathcal{L}}(\varphi),$$

then $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S .

- (b) Let \mathcal{L} be a logic such that $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S . Define the logic \mathcal{L}' by

$$\mathcal{L}'[\tau]_0 = \mathcal{L}[\tau]_0 \quad \text{and} \quad (\mathcal{A} \models^{\mathcal{L}'} \varphi \text{ iff } (\mathcal{A} \models^{\mathcal{L}} \varphi \text{ and } \mathcal{A} \in S[\tau])).$$

Then $\mathcal{L}' \equiv_{\text{es}} \mathcal{C}$ on S , too; moreover, \mathcal{L}' satisfies (2*).

The proof is left as an exercise. □

To get a first example we restate a result of Chapter 7:

Proposition 11.2.3 (a) $\text{FO}(\text{DTC}) \equiv_{\text{es}} \text{LOGSPACE}$ on the class of ordered structures.

⁷ Readers familiar with set theory will apologize our use of classes of classes.

(b) $\text{FO}(\text{IFP}) \equiv_{\text{es}} \text{PTIME}$ on the class of ordered structures.

Proof. (a) Let $\prec \in \tau$ and let $\mathcal{O}[\tau]$ be the class of ordered τ -structures (\prec being the ordering). Since $\mathcal{O}[\tau]_{\prec} \in \text{LOGSPACE}$ and $\text{FO}(\text{DTC})$ satisfies the effectivity conditions (1), (3) of Definition 11.2.1, it suffices to show that $\text{FO}(\text{DTC})$ satisfies (2*) of part (a) of the preceding proposition. Clearly, for $K \subseteq \mathcal{O}[\tau]$ we have (arguing as for Proposition 7.5.13),

$$K \in \text{LOGSPACE} \quad \text{iff} \quad K_{\prec} \in \text{LOGSPACE}$$

and therefore,

$$K_{\prec} \in \text{LOGSPACE} \quad \text{iff} \quad \text{there is an } \text{FO}(\text{DTC})[\tau]\text{-sentence } \varphi \\ \text{such that } K = \text{Mod}(\varphi).$$

The proof of (b) is similar. □

We now turn to the class **GRAPH** of finite graphs and prove that the main question is equivalent to its restriction to **GRAPH**. The reason is that any structure can feasibly be coded (interpreted) into a graph. We give the exact notion of interpretation.

Fix vocabularies σ and τ where $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i .⁸ An $(s+1)$ -tuple

$$\Pi := (\pi_{\text{uni}}(\bar{x}), \pi_{R_1}(\bar{x}_1, \dots, \bar{x}_{r_1}), \dots, \pi_{R_s}(\bar{x}_1, \dots, \bar{x}_{r_s}))$$

of $\mathcal{L}[\tau]$ -formulas (with free variables among the displayed ones), where for some $k \geq 1$ all tuples $\bar{x}, \bar{x}_1, \dots$ have length k , is called an \mathcal{L} -interpretation of σ in τ of width k .

Given a τ -structure \mathcal{A} , an \mathcal{L} -interpretation Π of σ in τ induces a σ -structure \mathcal{A}^{Π} over k -tuples of A , provided $\pi_{\text{uni}}^{\mathcal{A}}(-)$ is nonempty; the case $\pi_{\text{uni}}^{\mathcal{A}}(-) = \emptyset$ plays a marginal role for our considerations and will be neglected. The structure \mathcal{A}^{Π} has universe $\pi_{\text{uni}}^{\mathcal{A}}(-)$, the interpretation of R_i is given by

$$\pi_{R_i}^{\mathcal{A}}(-, \dots, -) \cap (\pi_{\text{uni}}^{\mathcal{A}}(-))^{r_i}.$$

For better reading we often suppress the intersection with the corresponding power of the universe and succinctly write

$$\mathcal{A}^{\Pi} = (\pi_{\text{uni}}^{\mathcal{A}}(-), \pi_{R_1}^{\mathcal{A}}(-, \dots, -), \dots, \pi_{R_s}^{\mathcal{A}}(-, \dots, -)).$$

Exercise 11.2.4 Let \mathcal{L} be one of the logics FO , $\text{FO}(\text{DTC})$, $\text{FO}(\text{TC})$, $\text{FO}(\text{IFP})$, $\text{FO}(\text{PFP})$, $L_{\infty\omega}^{\omega}$, or SO and let Π be an \mathcal{L} -interpretation of σ in τ . Show that for every $\mathcal{L}[\sigma]$ -sentence ψ there is an $\mathcal{L}[\tau]$ -sentence $\psi^{-\Pi}$ such that for all τ -structures \mathcal{A} (with $\pi_{\text{uni}}^{\mathcal{A}}(-) \neq \emptyset$),

$$\mathcal{A} \models \psi^{-\Pi} \quad \text{iff} \quad \mathcal{A}^{\Pi} \models \psi.$$

Check whether your proof works for **MSO**, too. If “yes”, it is incorrect. □

⁸ For simplicity, we restrict ourselves to relational σ . It should be clear how the notion of interpretation can be extended to vocabularies σ containing constants.

With the following proposition we show that “up to FO-interpretations, structures can be viewed as graphs”.

Proposition 11.2.5 *Let τ be a vocabulary. There are FO-interpretations Π of $\{E\}$ in τ and Δ of τ in $\{E\}$ with the following properties:*

(i) *For all τ -structures \mathcal{A} with at least two elements, $\mathcal{A}^\Pi \in \text{GRAPH}$ and*

$$(\mathcal{A}^\Pi)^\Delta \cong \mathcal{A}.$$

(ii) *For all $K \subseteq \text{Str}[\tau]$,*

$$K_{<} \in \text{PTIME} \quad \text{iff} \quad (K^\Pi)_{<} \in \text{PTIME},$$


where $K^\Pi := \{\mathcal{B} \mid \mathcal{B} \cong \mathcal{A}^\Pi \text{ for some } \mathcal{A} \in K\}$.

Proof. We treat the case $\tau = \{T, R\}$ with unary T and binary R . From the exposition it will be clear how to handle relation symbols of higher arity.

To get an idea of how to define Π , we first code a given τ -structure \mathcal{A} as a graph \mathcal{G} as follows:

- Each $a \in A$ is a vertex in \mathcal{G} ;
- each $a \in A$ has two new vertices a_1, a_2 linked to it;
- each $a \in T^A$ is linked to each element of a clique of two new elements;
- for each $(a, b) \in R^A$, a is linked to one element and b to the remaining two elements of a clique of three new elements.

The following picture shows \mathcal{G} in case $A = \{a, b\}$, $T^A = \{a\}$, and $R^A = \{(a, b)\}$.



Obviously, \mathcal{G} contains all the data of \mathcal{A} . As \mathcal{G} has more elements than \mathcal{A} , any interpretation Π such that $\mathcal{A}^\Pi \cong \mathcal{G}$ has to have a width $k \geq 2$. Then $(a, \dots, a) \in A^k$ can be the vertex of \mathcal{G} corresponding to $a \in A$. But, what elements of A^k may serve as “legs” a_1 and a_2 attached to a ? In general, for arbitrary \mathcal{A} there is no way to uniquely define two such elements for every $a \in A$. Thus, we replace each leg of a as well as every other auxiliary point attached to a via T or R by all elements of the set $A_a := \{b \in A \mid b \neq a\}$

(as A has at least two elements, A_a is nonempty). Altogether, we are led to the following modification \mathcal{H} of \mathcal{G} . We set $k = 9$ and let H contain, for every $a, b \in A$, the following points:

- (1) the point a^9 ($:= (a, \dots, a)$);
- (2) the points of $a^2 A_a a^6$ ($:= \{(a, a, c, a, \dots, a) \mid c \in A, c \neq a\}$) and of $a^3 A_a a^5$;
- (3) if $T^A a$, then the points of $a^4 A_a a^4$ and of $a^5 A_a a^3$;
- (4) if $R^A ab$, then the points of $aba^4 A_a a^2$, of $aba^5 A_a a$, and of $aba^6 A_a$.

Define the edge relation E^H as follows:

- for $a \in A$, connect the points in (2) with a^9
- for $a \in A$ with $T^A a$ and for $c \in A_a$, link $a^4 c a^4$ and $a^5 c a^3$ to a^9 and to each other
- for $a, b \in A$ with $R^A ab$ and for $c \in A_a$, form a clique out of $aba^4 c a^2$, $aba^5 c a$, and $aba^6 c$, and link $aba^4 c a^2$ to a^9 and $aba^5 c a$ and $aba^6 c$ to b^9 .

Let $\Pi = (\pi_{\text{uni}}(\bar{x}), \pi_E(\bar{x}, \bar{y}))$ with \bar{x}, \bar{y} of length 9 be the interpretation of $\{E\}$ in τ that describes this encoding. For instance, set

$$\begin{aligned} \pi_{\text{uni}}(\bar{x}) &:= x_1 = x_2 = \dots = x_9 \\ &\quad \vee (x_1 = x_2 = x_4 = \dots = x_9 \wedge x_3 \neq x_1) \\ &\quad \vee (x_1 = x_2 = x_3 = x_5 = \dots = x_9 \wedge x_4 \neq x_1) \\ &\quad \vee (T x_1 \wedge x_1 = \dots = x_4 = x_6 = \dots = x_9 \wedge x_5 \neq x_1) \\ &\quad \vee (T x_1 \wedge x_1 = \dots = x_5 = x_7 = x_8 = x_9 \wedge x_6 \neq x_1) \\ &\quad \vee \dots \end{aligned}$$

Then we have $\mathcal{A}^\Pi \in \text{GRAPH}$.

We now define an FO-interpretation Δ of τ in $\{E\}$ of width 1 that allows to regain the structure \mathcal{A} in \mathcal{A}^Π , more precisely: $(\mathcal{A}^\Pi)^\Delta \cong \mathcal{A}$ for all τ -structures \mathcal{A} with at least two elements.

In view of the definition of Π this is achieved by setting

$$\begin{aligned} \delta_{\text{uni}}(x) &:= \text{“}x \text{ has (at least) two legs” (}y \text{ is a leg of } x, \\ &\quad \text{if } y \text{ is only linked to } x) \\ \delta_T(x) &:= \text{“there are two elements which are linked to } x \text{ and} \\ &\quad \text{to each other and to no other point”} \\ \delta_R(x, y) &:= \text{“there are three points } u, v, w \text{ such that } u \text{ is linked to } x \\ &\quad \text{and } v, w \text{ are linked to } y, \text{ to } u, \text{ and to each other,} \\ &\quad \text{and } u, v, w \text{ are not linked to other points”}. \end{aligned}$$

We now come to part (ii). In the argumentation we tacitly assume that the τ -structures in question have at least two elements. (The remaining finitely many cases have to be treated separately.) First, we show that $(\text{Str}[\tau]^\Pi)_<$ is in PTIME. Namely, given an $\{E\}$ -structure \mathcal{B} , in polynomial time we can check whether

- \mathcal{B} is a graph
- all points with legs (“domain points”) have the same number of legs, this number being equal to twice the number of domain points minus two
- all other points are linked to exactly one domain point and the set of these points correctly codes relations T and R .

Now, let $K \subseteq \text{Str}[\tau]$. Assume that $K_{<} \in \text{PTIME}$. Given an ordered $\{E\}$ -structure $(\mathcal{B}, <)$, check whether $(\mathcal{B}, <) \in (\text{Str}[\tau]^\Pi)_{<}$. If not, $(\mathcal{B}, <) \notin (K^\Pi)_{<}$. Otherwise, let $\mathcal{B} \cong \mathcal{A}^\Pi$ for some τ -structure \mathcal{A} . Using Δ , construct the τ -structure \mathcal{B}^Δ in time polynomial in $\|\mathcal{B}\|$. By part (i) of our proposition, $\mathcal{B}^\Delta \cong (\mathcal{A}^\Pi)^\Delta \cong \mathcal{A}$. Let $<^\Delta$ be the ordering induced on \mathcal{B}^Δ by $<$. Then,

$$(\mathcal{B}, <) \in (K^\Pi)_{<} \quad \text{iff} \quad (\mathcal{B}^\Delta, <^\Delta) \in K_{<}$$

In fact, if $(\mathcal{B}^\Delta, <^\Delta) \in K_{<}$ then $\mathcal{A} \in K$ and hence, $(\mathcal{A}^\Pi \cong) \mathcal{B} \in K^\Pi$ and $(\mathcal{B}, <) \in (K^\Pi)_{<}$. Conversely, if $(\mathcal{B}, <) \in (K^\Pi)_{<}$ then $\mathcal{B} \cong \mathcal{A}_1^\Pi$ for some $\mathcal{A}_1 \in K$ and hence, $\mathcal{B}^\Delta \cong \mathcal{A}_1$ and $(\mathcal{B}^\Delta, <^\Delta) \in K_{<}$.

By assumption, the right hand side can be checked in time polynomial in $\|\mathcal{B}^\Delta\|$ and hence, in time polynomial in $\|\mathcal{B}\|$. Altogether, the left hand side can be checked in time polynomial in $\|\mathcal{B}\|$.

For the other direction assume that $(K^\Pi)_{<} \in \text{PTIME}$ and let $(\mathcal{A}, <)$ be an ordered $\tau_{<}$ -structure. We can construct \mathcal{A}^Π and, restricting the $<$ -lexicographic ordering of tuples to \mathcal{A}^Π , an ordering $<^\Pi$ on \mathcal{A}^Π in time polynomial in $\|\mathcal{A}\|$. Since, by (i),

$$(\mathcal{A}, <) \in K_{<} \quad \text{iff} \quad (\mathcal{A}^\Pi, <^\Pi) \in (K^\Pi)_{<}$$

and $(K^\Pi)_{<} \in \text{PTIME}$ by assumption, we altogether can decide the left hand side in time polynomial in $\|\mathcal{A}\|$. □

Now we are able to reduce the problem whether there is a logic \mathcal{L} with $\mathcal{L} \equiv_{\text{es}} \text{PTIME}$ to the class of graphs.

Theorem 11.2.6 *If there is a logic that effectively strongly captures PTIME on GRAPH then there is a logic effectively strongly capturing PTIME.*

Proof. Choose a logic \mathcal{L}_G with $\mathcal{L}_G \equiv_{\text{es}} \text{PTIME}$ on GRAPH. Define the logic \mathcal{L} as follows. For every vocabulary τ let $\mathcal{L}[\tau] = \mathcal{L}[\tau]_0$ be the set of sentences of $\mathcal{L}_G[\{E\}]$. For a τ -structure \mathcal{A} and $\varphi \in \mathcal{L}[\tau]$ define

$$(*) \quad \mathcal{A} \models^{\mathcal{L}} \varphi \quad \text{iff} \quad \mathcal{A}^\Pi \models^{\mathcal{L}_G} \varphi$$

(here Π is the interpretation of $\{E\}$ in τ as given by Proposition 11.2.5). We claim that $\mathcal{L} \equiv_{\text{es}} \text{PTIME}$.

Firstly, \mathcal{L} strongly captures PTIME: Assume $K \subseteq \text{Str}[\tau]$. We have to show that

$$K_{<} \in \text{PTIME} \quad \text{iff} \quad \text{for some } \varphi \in \mathcal{L}[\tau], K = \text{Mod}^{\mathcal{L}}(\varphi)$$

or, by Proposition 11.2.5(ii), that

$$(K^\Pi)_< \in \text{PTIME} \quad \text{iff} \quad \text{for some } \varphi \in \mathcal{L}[\tau], K = \text{Mod}^{\mathcal{L}}(\varphi).$$

If $(K^\Pi)_< \in \text{PTIME}$, then $K^\Pi = \text{Mod}^{\mathcal{L}_G}(\varphi)$ for some $\varphi \in \mathcal{L}_G[\{E\}]$, since $\mathcal{L}_G \equiv_{\text{es}} \text{PTIME}$. Therefore, by (*), $K = \text{Mod}^{\mathcal{L}}(\varphi)$. Conversely, assume $K = \text{Mod}^{\mathcal{L}}(\varphi)$ for some $\varphi \in \mathcal{L}[\tau]$. Then, $K^\Pi = \text{Mod}^{\mathcal{L}_G}(\varphi) \cap \text{Str}[\tau]^\Pi$ and therefore, $(K^\Pi)_< = \text{Mod}^{\mathcal{L}_G}(\varphi)_< \cap (\text{Str}[\tau]^\Pi)_<$. Both classes on the right hand side are in PTIME (the first one, as $\mathcal{L}_G \equiv_{\text{es}} \text{PTIME}$, the second one again by 11.2.5(ii)). Thus, $(K^\Pi)_< \in \text{PTIME}$.

Let τ be given. Of course, $\mathcal{L}[\tau]$ is recursive. Moreover, since $\mathcal{L}_G \equiv_{\text{es}} \text{PTIME}$, we effectively can assign to $\varphi \in \mathcal{L}[\tau]$ a pair (M_φ, d) , where M_φ is a deterministic x^d time-bounded machine accepting $\text{Mod}^{\mathcal{L}_G}(\varphi)_<$. To check whether a structure $(\mathcal{A}, <) \in \text{Str}(\tau)_<$ is in $\text{Mod}^{\mathcal{L}}(\varphi)_<$, by a polynomial time algorithm we pass to $(\mathcal{A}^\Pi, <^\Pi)$ (where $<^\Pi$ is the lexicographic ordering on A^Π induced by $<$) and then use M_φ to decide whether $(\mathcal{A}^\Pi, <^\Pi) \in \text{Mod}^{\mathcal{L}_G}(\varphi)_<$, and hence, by (*), whether $(\mathcal{A}, <) \in \text{Mod}^{\mathcal{L}}(\varphi)_<$. Clearly, we can effectively calculate a time bound.

Altogether, we have $\mathcal{L} \equiv_{\text{es}} \text{PTIME}$. □

Further notions and results of the preceding section can be restricted to classes of structures, too. E.g., let τ be a vocabulary and S a class of τ -structures and \mathcal{C} a deterministic complexity class. A map $C_\tau : \text{Str}[\tau] \rightarrow \text{Str}[\tau_<]$ is a \mathcal{C} -canonization on S , if

- (1) For all $\mathcal{A} \in S$, $\mathcal{A} \cong C_\tau(\mathcal{A})|_\tau$ and $<^{C_\tau(\mathcal{A})}$ is an ordering.
- (2) For all $\mathcal{A}, \mathcal{B} \in \text{Str}[\tau]$, if $\mathcal{A} \cong \mathcal{B}$ then $C_\tau(\mathcal{A}) \cong C_\tau(\mathcal{B})$.
- (3) C_τ is \mathcal{C} -computable, more precisely: there is a PTIME-algorithm that, applied to $(\mathcal{A}, <^{\mathcal{A}}) \in \text{Str}[\tau]_<$, gives (the encoding of an ordered version of) $C_\tau(\mathcal{A})$.

Analogously, one defines the notion of \mathcal{C} -invariantization. The proof of 11.1.8 showing that PTIME-invariantizations yield PTIME-canonizations, in general, does not work in the restricted case as it uses PTIME-invariantizations for structures with a larger vocabulary. However, in the case of GRAPH one can settle this difficulty by interpreting the expanded structures in graphs:

Exercise 11.2.7 (a) If there is a PTIME-canonization on GRAPH, then there is a PTIME-canonization.

(b) If there is a PTIME-invariantization on GRAPH, then there is a PTIME-invariantization. □

Exercise 11.2.8 Let τ be a vocabulary and S a class of τ -structures. Show: If there is a PTIME-canonization on S then there is a logic \mathcal{L} such that $\mathcal{L} \equiv_{\text{es}} \text{PTIME}$ on S . □

We now come to some positive results and start with a simple observation. We saw that the capturing result for PTIME proven in Chapter 7 can be reformulated as

FO(IFP) effectively strongly captures PTIME on the class of ordered structures.

This can easily be generalized to

Proposition 11.2.9 *Let $S \subseteq \text{Str}[\tau]$. Assume that there is an FO(IFP)[τ]-formula $\varphi(x, y, \bar{z})$ with free variables among x, y, \bar{z} such that for all $A \in S$ there are $\bar{b} \in A$ such that $\varphi^A(-, -, \bar{b})$ is an ordering of A . Then, $\text{FO(IFP)} \equiv_{\text{es}} \text{PTIME}$ on S .*

Proof. We first treat the case without parameters. The idea is to use the result just quoted, letting $\varphi(x, y)$ play the role of $x < y$. More explicitly: If $K_{<} \in \text{PTIME}$ then $K_{<} = \text{Mod}(\psi)$ for some sentence $\psi \in \text{FO(IFP)}[\tau_{<}]$. But then, for the FO(IFP)[τ]-sentence $\psi \frac{\varphi(\dots, \dots)}{\dots < \dots}$ we have that $K \cap S = \text{Mod}(\psi \frac{\varphi(\dots, \dots)}{\dots < \dots}) \cap S$. Conversely, if φ is an FO(IFP)[τ]-sentence and $K := \text{Mod}(\varphi)$, then we have $K_{<} \in \text{PTIME}$ and trivially, $K \cap S = \text{Mod}(\varphi) \cap S$.

For the general case one replaces the formula $\psi \frac{\varphi(\dots, \dots)}{\dots < \dots}$ by the formula $\exists \bar{y}(\varphi(-, -, \bar{y}) \text{ is an ordering} \wedge \psi \frac{\varphi(\dots, \dots, \bar{y})}{\dots < \dots})$. □

The following two exercises contain some applications and generalizations of the preceding proposition.

Exercise 11.2.10 (a) Show that $\text{FO(IFP)} \equiv_{\text{es}} \text{PTIME}$ on the class of cycles. Remember that a cycle is a graph that is isomorphic to some \mathcal{G}_l (cf. 2.3.8). Hint: Note that there is a formula $\varphi(x, y, u, v)$ in $\text{FO(IFP)}[\{E\}]$ such that for any cycle \mathcal{G} and all $a, b \in G$ with $E^{\mathcal{G}}ab$, the relation $\varphi^{\mathcal{G}}(-, -, a, b)$ is an ordering of G .

(b) A finite grid is a graph isomorphic to

$$(\{0, \dots, n\} \times \{0, \dots, m\}, \{((i, j), (k, l)) \mid i, k \leq n, j, l \leq m, |i - k| + |j - l| = 1\})$$

for some n, m . Show that $\text{FO(IFP)} \equiv_{\text{es}} \text{PTIME}$ on the class of grids. □


Note that in the situation given by the hypotheses of Proposition 11.2.9 we have, using the notations of its proof, that $C_{\tau} : \text{Str}[\tau] \rightarrow \text{Str}[\tau_{<}]$ with

$$C_{\tau}(\mathcal{A}) := (\mathcal{A}, \varphi^{\mathcal{A}}(-, -))$$

is a PTIME-canonicalization on S . Moreover, $C_{\tau}(\mathcal{A}) = \mathcal{A}^{\Pi}$ for the FO(IFP)-interpretation Π of $\tau_{<}$ in τ given by $\pi_{\text{uni}}(x) := x = x$, $\pi_{<}(x, y) := \varphi(x, y)$, and $\pi_R(\bar{x}) := R\bar{x}$ for $R \in \tau$. The proof of 11.2.9 contains the main argument needed to show the following exercise:

Exercise 11.2.11 Let \mathcal{L} be one of the logics mentioned in 11.2.4 and \mathcal{C} a deterministic complexity class with $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on ordered structures. Assume that S is a class of τ -structures and C_τ a \mathcal{C} -canonization on S given by an \mathcal{L} -interpretation of $\tau_{<}$ in τ (in the form $C_\tau : \mathcal{A} \mapsto \mathcal{A}^\Pi$). Then $\mathcal{L} \equiv_{\text{es}} \mathcal{C}$ on S . \square

Example 11.2.12 Let K_0 be the class of $\tau := \{E\}$ -structures isomorphic to some \mathcal{H}_l with $l \geq 3$, where \mathcal{H}_l is a digraph as pictured by




There is no formula that defines an ordering on all the structures in K_0 , even if parameters are allowed. In fact, consider a formula $\varphi(x, y, z_1, \dots, z_n)$ and the digraph \mathcal{H}_l for $l > n + 1$. Then, for any choice of parameters d_1, \dots, d_n , there are corresponding elements b, c of in-degree 1 (cf. the figure) not occurring among d_1, \dots, d_n . There is an automorphism of \mathcal{H}_l that fixes d_1, \dots, d_n and interchanges b and c . Hence,

$$(b, c) \in \varphi^{\mathcal{H}_l}(-, -, \bar{d}) \text{ iff } (c, b) \in \varphi^{\mathcal{H}_l}(-, -, \bar{d});$$

so $\varphi^{\mathcal{H}_l}(-, -, \bar{d})$ is not an ordering on \mathcal{H}_l .

Thus, we cannot apply 11.2.9 to conclude $\text{FO(IFP)} \equiv_{\text{es}} \text{PTIME}$ on K_0 . We obtain this result from the preceding exercise showing that there is a “definable” LOGSPACE-canonization. In fact, there is an FO(DTC) -interpretation $\Pi := (\pi_{\text{uni}}, \pi_E, \pi_{<})$ of $\sigma := \{E, <\}$ in $\tau (= \{E\})$ assigning to the digraph \mathcal{H}_l the ordered digraph \mathcal{H}_l^Π consisting of the isomorphic copy of \mathcal{H}_l



together with the ordering indicated by

$$(a_1, a_1) < (a_1, a_2) < (a_1, a_3) < (a_2, a_1) < (a_2, a_2) < \dots < (a_l, a_1).$$

To define Π , for $i = 1, 2, 3$ let $\varphi_i(x)$ be a formula defining the i -th “cross-point” of \mathcal{H}_l ,

$$\begin{aligned}
\varphi_1(x) &:= \forall y \neg Eyx \\
\varphi_2(x) &:= \exists z (\varphi_1(z) \wedge \exists y (Ezy \wedge Eyx)) \\
\varphi_3(x) &:= \exists z (\varphi_2(z) \wedge \exists y (Ezy \wedge Eyx)).
\end{aligned}$$

Then, for example, set

$$\pi_{\text{uni}}(x, y) := (\exists^2 z Exz \wedge (\varphi_1(y) \vee \varphi_2(y) \vee \varphi_3(y)) \vee (\forall z \neg Exz \wedge \varphi_1(y)).$$

Using the preceding exercise we conclude that $\text{FO}(\text{IFP}) \equiv_{\text{es}} \text{PTIME}$ on K_0 and that $\text{FO}(\text{DTC}) \equiv_{\text{es}} \text{LOGSPACE}$ on K_0 . \square

We turn to the class of trees. For this purpose we first state and prove a variation of the preceding exercise for fixed-point logic with counting, $\text{FO}(\text{IFP}, \#)$ (cf. subsection 8.4.1).

Proposition 11.2.13 *Let $\tau = \{R_1, \dots, R_m\}$ with r_i -ary R_i and let S be a class of τ -structures. Assume that there are $\text{FO}(\text{IFP}, \#)[\tau]$ -formulas $\varphi_1(\bar{\mu}_1), \dots, \varphi_m(\bar{\mu}_m)$ (where $\bar{\mu}_i$ is a sequence of number variables of length r_i) such that for every $A \in S$,*

$$(*) \quad A \cong (\{0, \dots, \|A\| - 1\}, \varphi_1^A(-), \dots, \varphi_m^A(-))^9$$

Then,

$$\text{FO}(\text{IFP}, \#) \equiv_{\text{es}} \text{PTIME on } S.$$

Proof. First, we observe that the map $C_\tau : \text{Str}[\tau] \rightarrow \text{Str}[\tau_{<}]$ given by

$$C_\tau(A) := (\{0, \dots, \|A\| - 1\}, \varphi_1^A(-), \dots, \varphi_m^A(-), <)$$

(with $<$ the natural ordering on $\{0, \dots, \|A\| - 1\}$) is a PTIME-canonization on S .

Let ψ be an $\text{FO}(\text{IFP})[\tau_{<}]$ -sentence. For every point variable x we choose a number variable μ_x (with $\mu_x \neq \mu_y$ for $x \neq y$), and we let ψ^* be the $\text{FO}(\text{IFP}, \#)[\tau]$ -sentence that arises from ψ by replacing every subformula $R_i x_1 \dots x_{r_i}$ by $\varphi_i(\mu_{x_1}, \dots, \mu_{x_{r_i}})$, by replacing every variable x by μ_x in all other atomic formulas, and by changing every subformula $\forall x \dots$ to $\forall \mu_x (\mu_x < \max \rightarrow \dots)$ and every subformula $\exists x \dots$ to $\exists \mu_x (\mu_x < \max \wedge \dots)$. Then we have for every τ -structure A :

$$A \models \psi^* \quad \text{iff} \quad C_\tau(A) \models \psi.$$

It remains to show that

$$\{K \cap S \mid K \subseteq \text{Str}[\tau], K_{<} \in \text{PTIME}\} = \{\text{Mod}(\varphi) \cap S \mid \varphi \in \text{FO}(\text{IFP}, \#)\}.$$

⁹ Note that $\{0, \dots, \|A\|\}$ is the universe of the number part of A ; thus, instead of $\varphi_i^A(-)$ we should write $\varphi_i^A(-) \cap \{0, \dots, \|A\| - 1\}^{r_i}$.

If φ is an FO(IFP, #)[τ]-sentence and $K = \text{Mod}(\varphi)$, then $K_{<} \in \text{PTIME}$ and clearly, $K \cap S = \text{Mod}(\varphi) \cap S$. Conversely, assume $K \subseteq \text{Str}[\tau]$ and $K_{<} \in \text{PTIME}$. Then $K_{<} = \text{Mod}(\psi)$ for some FO(IFP)[$\tau_{<}$]-sentence ψ . We claim that $K \cap S = \text{Mod}(\psi^*) \cap S$. In fact, for $\mathcal{A} \in S$,

$$\begin{aligned} \mathcal{A} \in K & \text{ iff } C_\tau(\mathcal{A}) \in K_{<} \quad (\text{as } \mathcal{A} \cong C_\tau(\mathcal{A})|\tau) \\ & \text{ iff } C_\tau(\mathcal{A}) \models \psi \\ & \text{ iff } \mathcal{A} \models \psi^* \quad (\text{see above}). \end{aligned}$$

□

We apply the preceding result to the class TREE of finite trees (cf. Chapter 1.A1).

Theorem 11.2.14 FO(IFP, #) \equiv_{es} PTIME on TREE.

In view of the preceding proposition it suffices to show:

Lemma 11.2.15 *There is a formula $\varphi(\mu, \nu)$ in FO(IFP)[$\{E\}$] such that for all $\mathcal{A} \in \text{TREE}$,*

$$\mathcal{A} \cong (\{0, \dots, \|\mathcal{A}\| - 1\}, \varphi^{\mathcal{A}}(-, -)).$$

Proof. Let $\mathcal{A} \in \text{TREE}$. By induction going from the leafs to the root, for each $u \in A$ we define a copy of the subtree \mathcal{A}_u of \mathcal{A} with root u on (an initial segment of) the number part of \mathcal{A} ; the tree relation of the copy of \mathcal{A}_u will be given by Xu_{--} (X is a ternary mixed relation variable with first component for point variables, the last two components for number variables). The case of leafs $u \in A$ being trivial, let $u \in A$ have the $E^{\mathcal{A}}$ -successors v_1, \dots, v_l . By induction hypothesis, the trees $\mathcal{A}_{v_1}, \dots, \mathcal{A}_{v_l}$ have isomorphic copies on initial segments of the number part of \mathcal{A} given by $Xv_{1--}, \dots, Xv_{l--}$. Using the ordering $<$ of the number part, these copies can be ordered lexicographically in an FO-definable way by

$$\begin{aligned} Xv_{i--} < Xv_{j--} \quad (\text{for short : } v_i < v_j) & \text{ iff } "Xv_{i--} \neq Xv_{j--}, \\ & \text{ and the minimal pair } (\mu, \nu), \text{ minimal in the lexicographic ordering} \\ & \text{ induced by } <, \text{ where they differ belongs to } Xv_{i--} ". \end{aligned}$$

We abbreviate $(\neg v_i < v_j \wedge \neg v_j < v_i)$ by $v_i \equiv v_j$. If, for example, $l = 4$ and $v_1 < v_2 \equiv v_3 < v_4$, we define the copy of \mathcal{A}_u on the number part of \mathcal{A} by assigning 0 to u and let it be followed by a copy of Xv_{1--} , two copies of Xv_{2--} , and one copy of Xv_{4--} . More precisely, in the general case we define

$$Xu\mu\nu$$

by the disjunction of the following formulas φ_1, φ_2 (φ_1 describes that the number 0 assigned to u is linked to the numbers assigned to the v_i ; φ_2 describes the copies associated with the subtrees \mathcal{A}_{v_i}).

Using $\delta(u, v)$ for the number term $\#\{w \mid \exists v'(Euv' \wedge v' < v \wedge w \in \mathcal{A}_{v'})\}$ which counts the number of all elements in trees $\mathcal{A}_{v'}$ of E -successors v' of u that are $<$ v , we set

$$\varphi_1(u, \mu, \nu, X) := \mu = \min \wedge \exists v(Euv \wedge \exists \zeta(0 \leq \zeta < \#\{w \mid Euv \wedge w \equiv v\} \\ \wedge \nu = 1 + \delta(u, v) + \zeta \cdot \#\{w \mid w \in \mathcal{A}_v\}));$$

$$\varphi_2(u, \mu, \nu, X) := \exists v(Euv \wedge \exists \zeta(0 \leq \zeta < \#\{w \mid Euv \wedge w \equiv v\} \wedge \\ \exists \xi \exists \eta(Xv\xi\eta \\ \wedge \mu = \xi + 1 + \delta(u, v) + \zeta \cdot \#\{w \mid w \in \mathcal{A}_v\} \\ \wedge \nu = \eta + 1 + \delta(u, v) + \zeta \cdot \#\{w \mid w \in \mathcal{A}_v\})).$$

Note that, with u_0 the root of \mathcal{A} ,

$Xu_0 _ _$ defines a copy of \mathcal{A} .

According to Lemma 7.3.11, addition and multiplication on the number part of \mathcal{A} are FO(IFP)-definable as is “ $w \in \mathcal{A}_v$ ”. Thus, the desired FO(IFP, #)[$\{E\}$]-formula $\chi(\mu, \nu)$, defining an isomorphic copy of \mathcal{A} on its number part, can be given as

$$\chi(\mu, \nu) := \exists u_0(u_0 \text{ is the root} \wedge [\text{IFP}_{u\mu\nu, X}(\varphi_1 \vee \varphi_2)]u_0\mu\nu). \quad \square$$

Exercise 11.2.16 (a) Show that $\text{FO}(\text{IFP}, \#) \equiv_{\text{es}} \text{PTIME}$ on the class of acyclic connected graphs. Hint: Each point of an acyclic connected graph can be viewed as the root of a corresponding tree.

(b) Show that $\text{FO}(\text{IFP}, \#) \equiv_{\text{es}} \text{PTIME}$ on the class of acyclic graphs. \square

In 11.2.6 we have seen that, in order to obtain a logic \mathcal{L} with $\mathcal{L} \equiv_{\text{es}} \text{PTIME}$, it is enough to find a logic \mathcal{L}' such that $\mathcal{L}' \equiv_{\text{es}} \text{PTIME}$ on GRAPH. We then have seen that there are subclasses of GRAPH such as the class of acyclic graphs or the class of trees that can effectively strongly be captured by a logic. Further positive results are contained in [65, 68] where it is shown, for example, that $\text{FO}(\text{IFP}, \#) \equiv_{\text{es}} \text{PTIME}$ on the class of planar graphs. These positive results may be interpreted as an approximation to a positive solution of the main question, if there should be a positive solution. If not, they might serve to discover the borderline where the possibility for an effective strong capturing gets lost.

Notes 11.2.17 The main question whether there is a logic strongly capturing PTIME was first asked in [19]. The notion of a logic effectively strongly capturing a complexity class originates with Gurevich [71]; Theorem 11.1.9 is from [73]. The capturing result 11.2.14 goes back to [97, 117]. For further capturing results see [82] and the survey paper [66].

12. Quantifiers and Logical Reductions

In Chapter 7 we have seen that the fixed-point logic FO(IFP) captures PTIME. In the last chapter we have remarked that the question

(*) *Is there a logic that effectively strongly captures PTIME?*

is still open. We have discussed strategies for a solution and thereby given some positive results in case we restrict ourselves to certain classes of structures.

The notion of logic underlying the main question is very general. For instance, we have admitted that the sentences of a logic are Turing machines (and the models of a Turing machine are the structures it accepts). Thus, its syntax lacks the characteristic flavour of logics we are accustomed to. A corresponding positive answer to (*) could hardly be considered to be a *logical* characterization of PTIME. However, in the present chapter we shall show that if there is a logic that effectively strongly captures PTIME, then there is a logic with a familiar syntax.

At a first glance the following idea could lead to such a logic: Consider, for example, FO(IFP). We know that there is some class K of structures such that $K_{<}$ belongs to PTIME, but K is not axiomatizable in FO(IFP). In a first step one could add a “logical construct” to FO(IFP) that allows for an axiomatization of K , and in this way one could continue trying to enlarge the expressive power just up to PTIME. In traditional model theory there is a well-established method to enrich logics by new concepts, namely by adjoining so-called Lindström quantifiers. We have already met such quantifiers, for example, the counting quantifiers (cf. section 3.4).

In the first section we shall introduce the concept of Lindström quantifier and demonstrate its scope by various examples showing that some important logics encountered so far can be obtained as extensions of first-order logic by such quantifiers. In section 2 we prove some limiting results saying us, for example, that no extension of FO(IFP) by a finite set of Lindström quantifiers will strongly capture PTIME. Nevertheless, we shall see there that, if any reasonable logic strongly captures PTIME, then there is already such a logic that can be obtained from first-order logic by adding a simple infinite set of quantifiers. When pursuing this line also for other complexity classes such as NPTIME, a logical analogue of complexity-theoretical reductions turns out to be useful. Section 3 presents results of the corresponding kind.

If we extend a logic \mathcal{L} by the binary Lindström quantifier Q where $Qxy\varphi(x, y)$ says that $\varphi(x, y)$ defines a connected graph on the universe, then in the extended logic $\mathcal{L}(Q)$ we can speak about the connectivity of definable graphs by simply using the quantifier Q . There is a similarity to Turing machines with oracles that have the ability to supply, in one step, correct answers to the question whether a binary relation (encoded on some tape) is a connected graph. We shall see in the last section that this similarity is by no means superficial.

Throughout the chapter all vocabularies are relational and structure means finite structure.

12.1 Lindström Quantifiers

For a binary relation symbol E we know that FO(TC), but not FO, allows to express

(*) (x, y) is in the transitive closure of E .

What do we expect from any “natural” logic \mathcal{L} allowing to formalize (*)? Clearly, for every formula $\varphi(u, v)$ of \mathcal{L} we want to have a formula in \mathcal{L} expressing

(x, y) is in the transitive closure of $\varphi(-, -)$.

Similarly, in a natural logic \mathcal{L} which allows to express that a graph has a Hamiltonian path, we expect to be able to say that an \mathcal{L} -definable graph has a Hamiltonian path. If in the following definitions we take as K the class of graphs with a Hamiltonian path, we get a quantifier Q such that FO(Q) will be the smallest natural extension of FO having this ability.

Fix a vocabulary

$$\sigma = \{R_1, \dots, R_s\}$$

where R_i is r_i -ary, and a class K of σ -structures. Furthermore, let \mathcal{L} be some logic such as FO, FO(IFP), or $L_{\infty\omega}$. For a new “quantifier symbol” Q_K we enlarge, for any vocabulary τ , the formation rules for $\mathcal{L}[\tau]$ -formulas by the following clause:

(+) If for $1 \leq i \leq s$, $\psi_i(\bar{x}_i)^1$ is a formula and the length of \bar{x}_i is r_i then

$$Q_K \bar{x}_1, \dots, \bar{x}_s [\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)]$$

is a formula, too.

¹ We only display the variables relevant in the given context. The variables in each \bar{x}_i are pairwise distinct, but the same variable may occur in distinct \bar{x}_i 's.

In case $s = 1$ we write $Q\bar{x}\psi(\bar{x})$ instead of $Q\bar{x}[\psi(\bar{x})]$. A variable y is free in $Q_K\bar{x}_1, \dots, \bar{x}_s[\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)]$ if, for some i , we have $y \in \text{free}(\psi_i)$ and y not in \bar{x}_i .

For any τ -structure \mathcal{A} the meaning of $Q_K\bar{x}_1, \dots, \bar{x}_s[\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)]$ is given by (we suppress the assignment to variables)

$$\mathcal{A} \models Q_K\bar{x}_1, \dots, \bar{x}_s[\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)] \quad \text{iff} \quad (A, \psi_1^{\mathcal{A}}(-), \dots, \psi_s^{\mathcal{A}}(-)) \in K$$

where $\psi_i^{\mathcal{A}}(-)$ stands for $\{\bar{b} \in A^{r_i} \mid \mathcal{A} \models \psi_i[\bar{b}]\}$.

Q_K is the *Lindström quantifier* (by short: *quantifier*) given by K . Its arity is defined to be $\max\{r_i \mid 1 \leq i \leq s\}$. In case $s = 1$ one speaks of a *simple* quantifier.

The extension of \mathcal{L} by Q_K is denoted by $\mathcal{L}(Q_K)$. If \mathbf{Q} is a class of Lindström quantifiers, the extension $\mathcal{L}(\mathbf{Q})$ of \mathcal{L} is defined by adding clause (+) for all $Q \in \mathbf{Q}$. $\mathcal{L}(\mathbf{Q})$ is called a *Lindström extension* of \mathcal{L} .

Obviously, we have for all σ -structures \mathcal{A} :

$$\mathcal{A} \models Q_K\bar{x}_1, \dots, \bar{x}_s[R_1\bar{x}_1, \dots, R_s\bar{x}_s] \quad \text{iff} \quad \mathcal{A} \in K.$$

Hence, K is axiomatizable in $\mathcal{L}(Q_K)$. Indeed, $\mathcal{L}(Q_K)$ is the least extension of \mathcal{L} that allows to axiomatize K . A proof requires some closure conditions of the logics involved. They are satisfied in the cases covered by the next exercise. As a rule, such closure conditions (cf. 12.1.11) will also be necessary in the following when dealing with statements on arbitrary logics and will be tacitly assumed.

Exercise 12.1.1 Let \mathcal{L}_1 be a Lindström extension of FO and \mathcal{L}_2 one of the logics FO(TC), FO(DTC), FO(IFP), FO(PFP), or $L_{\infty\omega}$. Show: If $\mathcal{L}_1 \leq \mathcal{L}_2$ and the class K of structures is axiomatizable in \mathcal{L}_2 then $\mathcal{L}_1(Q_K) \leq \mathcal{L}_2$. \square

One can show that any logic \mathcal{L} above FO that obeys the closure conditions referred to above, can be represented as a Lindström extension FO(\mathbf{Q}) of FO, where the quantifiers in \mathbf{Q} correspond to the classes axiomatizable in \mathcal{L} , i.e., $\mathcal{L} \equiv \text{FO}(\{Q_K \mid K \text{ axiomatizable in } \mathcal{L}\})$. Of course, one is interested in small sets \mathbf{Q} such that FO(\mathbf{Q}) and \mathcal{L} are equivalent.

We give some examples of Lindström quantifiers.

Examples 12.1.2 (a) Let $\sigma := \{P\}$ with unary P . For $K := \{(A, B) \mid \emptyset \neq B \subseteq A\}$, the simple unary quantifier Q_K corresponds to the existential quantifier, as $Q_K x\psi(x)$ is equivalent to $\exists x\psi(x)$.

(b) Similarly, for $K := \{(A, A) \mid A \text{ a nonempty set}\}$ we obtain the universal quantifier.

(c) For $n \geq 1$ and $K_n := \{(A, B) \mid B \subseteq A \text{ and } \|B\| \geq n\}$ the quantifier Q_{K_n} corresponds to the quantifier $\exists^{\geq n}$ (cf. section 3.4).

(d) For the class CONN of connected graphs we get a simple binary quantifier Q_{CONN} . So

$$Q_{\text{CONN}}xy \psi(x, y)$$

is valid in a structure \mathcal{A} iff $(\mathcal{A}, \psi^{\mathcal{A}}(-, -))$ is a connected graph.

(e) Let $\sigma := \{E, R\}$ with binary relation symbols E and R . If K is the class of σ -structures

$$\mathcal{A} = \{(A, E^A, R^A) \mid R^A \subseteq \text{TC}(E^A)\},$$

then Q_K can replace the unary transitive closure operator, as

$$[\text{TC}_{x,y}\psi(x, y)]st$$

is equivalent to

$$Q_Kxy, uv[\psi(x, y), u = s \wedge v = t],$$

and

$$Q_Kxy, uv[\psi_1(x, y), \psi_2(u, v)]$$

is equivalent to

$$\forall u \forall v (\psi_2(u, v) \rightarrow [\text{TC}_{x,y}\psi_1(x, y)]uv).$$

Thus, $\text{FO}(\text{TC}^1) \equiv \text{FO}(Q_K)$.

(f) Let K be a class of σ -structures, $<$ a new binary relation symbol, and \mathcal{L} a logic. Then,

$$\mathcal{L}(Q_{K<}) \equiv \mathcal{L}(Q_K) \text{ on ordered structures.}$$

In fact, if $\sigma = \{R_1, \dots, R_s\}$ and $\varphi_1, \dots, \varphi_s, \psi$ are formulas in some vocabulary τ that contains a binary relation symbol $<$, then on ordered τ -structures (that is, structures in which $<$ is an ordering) the formulas

$$Q_K\bar{x}_1, \dots, \bar{x}_s[\varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s)]$$

and

$$Q_{K<}\bar{x}_1, \dots, \bar{x}_s, xy[\varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s), x < y]$$

are equivalent, and the formulas

$$Q_{K<}\bar{x}_1, \dots, \bar{x}_s, xy[\varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s), \psi(x, y)]$$

and

$$\text{“}\psi(x, y) \text{ defines an ordering”} \wedge Q_K\bar{x}_1, \dots, \bar{x}_s[\varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s)]$$

are equivalent. Conclude: If $\mathcal{L}(\mathbf{Q})$ captures, say PTIME, then $K_{<} \in \text{PTIME}$ for all K such that $Q_K \in \mathbf{Q}$. \square

Exercise 12.1.3 Let K be a class such that $K_{<} \in \text{PTIME}$. Show that $\text{Mod}(\psi)_{<} \in \text{PTIME}$ for any sentence $\psi \in \text{FO}(\text{IFP})(Q_K)$. The same is true for sets of such quantifiers. Give the corresponding statements for PSPACE and FO(PFP) and for LOGSPACE and FO(DTC). Hint: For LOGSPACE cf. Appendix 7.5.1. \square

Finitely many Lindström quantifiers can be replaced by a single one. We give a precise statement in the next theorem. The following two technical remarks (1) and (2) will be useful for the proof. As above, let $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i and let K be a class of σ -structures.

(1) For a new r -ary relation symbol R let K_1 be the class of $\sigma \cup \{R\}$ -structures

$$K_1 := \{(A, R^A) \mid A \in K \text{ and } R^A = A^r\}.$$

Then $\mathcal{L}(Q_K) \equiv \mathcal{L}(Q_{K_1})$, since

$$\models_{\text{fin}} Q_K \bar{x}_1, \dots, \bar{x}_s [\psi_1, \dots, \psi_s] \leftrightarrow Q_{K_1} \bar{x}_1, \dots, \bar{x}_s, \bar{x} [\psi_1, \dots, \psi_s, \bar{x} = \bar{x}]$$

and

$$\models_{\text{fin}} Q_{K_1} \bar{x}_1, \dots, \bar{x}_s, \bar{x} [\psi_1, \dots, \psi_s, \psi] \leftrightarrow (Q_K \bar{x}_1, \dots, \bar{x}_s [\psi_1, \dots, \psi_s] \wedge \forall \bar{x} \psi(\bar{x})).$$

(2) Let $r \geq \max\{r_i \mid 1 \leq i \leq s\}$ and let σ' arise from σ by replacing each R_i by an r -ary relation symbol R'_i . Finally, let K' be the class of σ' -structures A' that we get from structures A in K by replacing each R_i^A by $R_i^{A'} := R_i^A \times A^{r-r_i}$. Using dummy variables, one easily shows that $\mathcal{L}(Q_K) \equiv \mathcal{L}(Q_{K'})$.

Theorem 12.1.4 *Let \mathcal{L} be a logic. Furthermore, let \mathbf{Q} be a finite set of Lindström quantifiers. Then there is a Lindström quantifier Q_K such that $\mathcal{L}(\mathbf{Q}) \equiv \mathcal{L}(Q_K)$. In particular, if the arities of the quantifiers in \mathbf{Q} are $\leq l$, then Q_K can be chosen of arity $l + 1$.*

Proof. Let \mathbf{Q} be $\{Q_{K_j} \mid 1 \leq j \leq k\}$ where K_j is a class of σ_j -structures. According to the preceding remarks we may assume that the σ_j have the same cardinality and that the Q_{K_j} are of the same arity l as are all relation symbols in the vocabularies σ_j . For simplicity of notation we treat the case of two simple binary quantifiers Q_{K_1} and Q_{K_2} . Let R be a ternary relation symbol. If no structure of the form (A, \emptyset) belongs to K_1 or K_2 , we define the class K of $\{R\}$ -structures by

$$\begin{aligned} K := & \{(A, B \times \{a\}) \mid (A, B) \in K_1, a \in A\} \\ & \cup \{(A, B \times \{a, b\}) \mid (A, B) \in K_2, a, b \in A, a \neq b\}. \end{aligned}$$

Then the following formulas are equivalent on structures with at least two elements (structures of cardinality 1 can be taken into consideration explicitly):

- (1) $Q_{K_1} xy\varphi(x, y)$ and $\exists v Q_K xyz(\varphi(x, y) \wedge z = v)$
- (2) $Q_{K_2} xy\varphi(x, y)$ and $\exists v \exists w (v \neq w \wedge Q_K xyz(\varphi(x, y) \wedge (z = v \vee z = w)))$
- (3) $Q_K xyz\psi(x, y, z)$ and $(\exists^=1 z \exists x \exists y \psi(x, y, z) \wedge Q_{K_1} xy \exists z \psi(x, y, z))$

$$\vee (\exists^=2 z \exists x \exists y \psi(x, y, z) \wedge \forall xy (\exists z \psi(x, y, z) \rightarrow \exists^=2 z \psi(x, y, z)))$$

$$\wedge Q_{K_2} xy \exists z \psi(x, y, z)).$$

If $K_1 \cup K_2$ contains structures of the form (A, \emptyset) , we add to K the structure $(A, A \times A \times \{a, b, c\})$ with distinct a, b, c if $(A, \emptyset) \in K_1$, and the structure $(A, A \times A \times \{a, b, c, d\})$ with distinct a, b, c, d if $(A, \emptyset) \in K_2$. Arguing similarly as above one shows that $\mathcal{L}(\{Q_{K_1}, Q_{K_2}\}) \equiv \mathcal{L}(Q_K)$. \square

Exercise 12.1.5 Generalize the preceding theorem to a countable set \mathbf{Q} of simple binary quantifiers, thereby weakening the assertion $\mathcal{L}(\mathbf{Q}) \equiv \mathcal{L}(Q_K)$ to $\mathcal{L}(\mathbf{Q}) \leq \mathcal{L}(Q_K)$. \square

$A \models [\text{TC}_{x,y}\varphi(x,y)]st$ means that (s, t) is in the transitive closure of the binary relation on A given by $\varphi^A(-, -)$. Similarly, for $k \geq 2$ and $\psi(x_1 \dots x_k, y_1 \dots y_k)$ we can read $A \models [\text{TC}_{\bar{x}, \bar{y}}\psi(\bar{x}, \bar{y})]s\bar{t}$ as “ (\bar{s}, \bar{t}) is in the transitive closure of the *binary* relation on A^k given by $\psi^A(-, -)$ ”. Guided by this example we define the k -vectorization of a Lindström quantifier.

Let $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i and let $k \geq 1$. Let $\sigma(k)$ arise from σ by replacing each relation symbol R in σ , say of arity r , by an $r \cdot k$ -ary relation symbol $R(k)$. An r -ary relation on A^k can also be viewed as an $r \cdot k$ -ary relation on A just by identifying the r -tuple consisting of the elements $\bar{a}_1, \dots, \bar{a}_r$ of A^k with the $r \cdot k$ -tuple $\bar{a}_1 \dots \bar{a}_r$ of elements of A . Using this identification on the level of tuples and relations, we define:

Definition 12.1.6 Let K be a class of σ -structures.

(a) K^k , the k -vectorization of K , is the class of $\sigma(k)$ -structures

$$K^k := \{(A, S_1, \dots, S_s) \mid (A^k, S_1, \dots, S_s) \in K\};$$

here, in (A, S_1, \dots, S_s) , the relation S_i is viewed as an $r_i \cdot k$ -ary relation over A , and in (A^k, S_1, \dots, S_s) it is viewed as an r_i -ary relation over A^k ; in (A, S_1, \dots, S_s) it is the interpretation of $R_i(k)$, in (A^k, S_1, \dots, S_s) the interpretation of R_i .

(b) We set $\mathbf{Q}_K^\omega := \{Q_{K^k} \mid k \geq 1\}$. \square

Thus, for the class K of Example 12.1.2(a) we have

$$\begin{aligned} A \models Q_{K^2}xy\varphi(x,y) & \text{ iff } (A^2, \varphi^A(-, -)) \in K \\ & \text{ iff } A \models \exists x\exists y\varphi(x,y), \end{aligned}$$

and for the class CONN (cf. 12.1.2(d))

$$A \models Q_{\text{CONN}^2}xx'yy'\varphi(xx', yy') \text{ iff } (A^2, \varphi^A(-, -)) \text{ is a connected graph.}$$

Exercise 12.1.7 Let K be as in Example 12.1.2(e) and let $k \geq 1$. Generalize the equivalence $\text{FO}(Q_K) \equiv \text{FO}(\text{TC}^1)$ to $\text{FO}(Q_{K^k}) \equiv \text{FO}(\text{TC}^k)$. Hence, $\text{FO}(\mathbf{Q}_K^\omega) \equiv \text{FO}(\text{TC})$. \square

In view of 9.6.1 we have for the class K of the preceding exercise:

$$K \text{ but not } K^2 \text{ is axiomatizable in } \text{FO}(\text{TC}^1).$$

The following exercise shows that our standard logics are closed under vectorizations.

Exercise 12.1.8 Let \mathcal{L} be one of the logics FO(TC), FO(DTC), FO(IFP), FO(PFP), or $L_{\infty\omega}$. Show: If K is axiomatizable in \mathcal{L} and $k \geq 1$, then K^k is axiomatizable in \mathcal{L} . Conclude: If \mathcal{L}_1 with $\mathcal{L}_1 \leq \mathcal{L}$ is a Lindström extension of FO and K is axiomatizable in \mathcal{L} , then $\mathcal{L}_1(Q_K^\omega) \leq \mathcal{L}$. \square

Two further logics which we have extensively studied can be obtained from FO by adding the vectorizations of a simple quantifier:

Theorem 12.1.9 *There are classes I and P of structures such that*

$$\text{FO(IFP)} \equiv \text{FO}(\mathbf{Q}_I^\omega) \quad \text{and} \quad \text{FO(PFP)} \equiv \text{FO}(\mathbf{Q}_P^\omega).$$

Proof. We first consider FO(IFP) and recall (cf. 9.4.3) that any FO(IFP)-formula is equivalent to a formula of the form

$$\forall u [\text{IFP}_{\bar{x}, X}(\psi \vee \exists \bar{y} \in X \exists \bar{z} \notin X \chi)] \tilde{u}$$

where ψ and χ are quantifier-free and X does not occur in ψ, χ . This normal form motivates the following definition of I.

Let $\sigma = \{U, R, V\}$ with ternary R and unary U, V and let I consist of the σ -structures \mathcal{A} such that there is a sequence A_0, A_1, A_2, \dots of subsets of A with the following properties:

- (1) $A_0 = \emptyset$ and $A_1 = U^{\mathcal{A}}$.
- (2) For $i \geq 1$, $A_{i+1} = A_i \cup \{a \in A \mid R^{\mathcal{A}}abc \text{ for some } b \in A_i \text{ and } c \notin A_i\}$.
- (3) There is an $i \geq 1$ such that $V^{\mathcal{A}} \subseteq A_i$.

To prove $\text{FO}(\mathbf{Q}_I^\omega) \leq \text{FO(IFP)}$, by 12.1.8, it suffices to show that I is axiomatizable in FO(IFP). In fact, I is the class of models of φ_0 where

$$\varphi_0 := \forall u (Vu \rightarrow [\text{IFP}_{x, X}(Ux \vee \exists y \in X \exists z \notin X Rxyz)] u).$$

For the other direction we consider an FO(IFP)-sentence φ in the normal form mentioned above, say

$$\varphi = \forall u [\text{IFP}_{\bar{x}, X}(\psi(\bar{x}) \vee \exists \bar{y} \in X \exists \bar{z} \notin X \chi(\bar{x}, \bar{y}, \bar{z}))] \tilde{u},$$

where X is k -ary. Given a structure \mathcal{A} , let X_0, X_1, X_2, \dots be the stages of the fixed-point operation in φ . Then $X_1 = \psi^{\mathcal{A}}(-)$, and \mathcal{A} is a model of φ iff $\{\tilde{a} \mid a \in A\} \subseteq X_\infty$. Hence, φ is equivalent to

$$Q_{I^k} \bar{x}, \bar{y}, \bar{z}, \bar{u} [\psi(\bar{x}), \chi(\bar{x}, \bar{y}, \bar{z}), u_1 = \dots = u_k].$$

The proof for FO(PFP) is similar; one uses the class P that is defined as I with (2), (3) modified to

- (2)' For $i \geq 1$, $A_{i+1} = \{a \in A \mid R^{\mathcal{A}}abc \text{ for some } b \in A_i \text{ and } c \notin A_i\}$.
- (3)' $V^{\mathcal{A}} = \emptyset$ or there is an $i \geq 1$ such that $V^{\mathcal{A}} \subseteq A_i$ and $A_i = A_{i+1}$. \square

We finish by representing FO(ATC) as a Lindström extension of FO. Setting $\sigma := \{E, U, R\}$ with binary E, R and unary U , we define \mathbf{A} to be the class of σ -structures (A, E^A, U^A, R^A) such that R^A is contained in the alternating transitive closure of (E^A, U^A) . Then one easily gets:

Theorem 12.1.10 $\text{FO(ATC)} \equiv \text{FO}(\mathbf{Q}_{\mathbf{A}}^{\omega})$. □

Notes 12.1.11 At many places of this chapter we tacitly assume that \mathcal{L} is a logic extending FO and satisfying certain closure conditions such as closure under negations or conjunctions (cf. (3) in the proof of Theorem 12.1.4). For instance, closure of \mathcal{L} under conjunctions means:

If φ and ψ are \mathcal{L} -formulas then $(\varphi \wedge \psi)$ is an \mathcal{L} -formula (or, at least, \mathcal{L} contains a formula equivalent to $(\varphi \wedge \psi)$).

Compare [31] for the precise statement and an analysis of such closure conditions.

12.2 PTIME and Quantifiers

In this section we shall see that no extension of FO(IFP) by means of finitely many Lindström quantifiers strongly captures PTIME. The same is true for all fragments of the extension of $\mathbf{L}_{\infty\omega}^{\omega}$ by all counting quantifiers.

For a set \mathbf{Q} of Lindström quantifiers, $\mathbf{L}_{\infty\omega}^{\omega}(\mathbf{Q})$ consists of those formulas of $\mathbf{L}_{\infty\omega}^{\omega}(\mathbf{Q})$ that contain only finitely many variables. As a first result we have

Theorem 12.2.1 *Let \mathbf{Q} be a finite set of Lindström quantifiers Q_L with $L_{<}$ in PTIME. Then for any relational vocabulary τ there is a class K of τ -structures such that $K_{<}$ belongs to PTIME, but K is not axiomatizable in $\mathbf{L}_{\infty\omega}^{\omega}(\mathbf{Q})$.*

The proof will be given below. As $\text{FO(IFP)}(\mathbf{Q}) \leq \mathbf{L}_{\infty\omega}^{\omega}(\mathbf{Q})$ (mimic the proof of $\text{FO(PFP)} \leq \mathbf{L}_{\infty\omega}^{\omega}$, cf. 8.4.2) we have

Corollary 12.2.2 *No extension of FO(IFP) by means of finitely many Lindström quantifiers strongly captures PTIME.* □

To give an application of 12.2.1, let E be binary and \mathbf{Q} be the set of all simple binary quantifiers corresponding to classes K of $\{E\}$ -structures with $K_{<}$ in PTIME. According to 12.1.5, \mathbf{Q} may be replaced by one (ternary) quantifier Q_L such that $\text{FO}(\mathbf{Q}) \leq \text{FO}(Q_L)$. Therefore, $\text{FO}(Q_L)$ and hence $\mathbf{L}_{\infty\omega}^{\omega}(Q_L)$, is able to axiomatize all the classes K of $\{E\}$ -structures with $K_{<}$ in PTIME. By the theorem, $L_{<}$ cannot be in PTIME, that is, $\text{FO}(Q_L)$ is strictly stronger than $\text{FO}(\mathbf{Q})$ (cf. 12.1.3).

The theorem has a generalization to certain infinite sets of quantifiers. Without proof we mention the following generalization to arity-bounded sets of quantifiers:

Theorem 12.2.3 *Let $n \geq 1$ and let \mathbf{Q} consist of Lindström quantifiers of arity $\leq n$. Then there is a relational vocabulary τ and a class K of τ -structures such that $K_{<}$ is in PTIME, but K is not axiomatizable in $L_{\infty\omega}^{\omega}(\mathbf{Q})$.*

Corollary 12.2.4 *No extension of FO(IFP) by means of Lindström quantifiers of bounded arity strongly captures PTIME. \square*

Since the first examples of PTIME classes not axiomatizable in FO(IFP) were related to the class EVEN of structures of even cardinality, the logic FO(IFP, #), fixed-point logic with counting, appeared as a natural candidate for a logic strongly capturing PTIME. However: Let \mathbf{C} be the set of counting quantifiers, i.e., $\mathbf{C} := \{Q_{K_n} \mid n \geq 1\}$, where K_n is as in 12.1.2(c). Then we have

Theorem 12.2.5 *There is a class K of graphs such that $K_{<}$ is in PTIME, but K is not axiomatizable in $L_{\infty\omega}^{\omega}(\mathbf{C})$.*

Note that $L_{\infty\omega}^{\omega}(\mathbf{C}) \equiv C_{\infty\omega}^{\omega}$. By 8.4.18, $\text{FO}(\text{IFP}, \#) \leq C_{\infty\omega}^{\omega}$. Therefore:

Corollary 12.2.6 *FO(IFP, #) does not strongly capture PTIME. \square*

In the following we prove Theorem 12.2.1 and Theorem 12.2.5.

Proof (of 12.2.1). First we show that for $\tau = \emptyset$ there is a class K of structures such that $K_{<}$ is in PTIME, but K is not axiomatizable in $L_{\infty\omega}^{\omega}(\mathbf{Q})$.

So let $\tau = \emptyset$ and fix $s \geq 1$. Thus τ -structures are sets and will be denoted by their universes. We start by proving that every $L_{\infty\omega}^{\omega}(\mathbf{Q})$ -sentence (of vocabulary $\tau = \emptyset$ and) with variables among v_1, \dots, v_s is equivalent to an FO(\mathbf{Q})-sentence of a certain simple form.

First we note that there are, up to logical equivalence, only finitely many quantifier-free formulas with variables among v_1, \dots, v_s and hence, only finitely many formulas $\psi_1^s, \dots, \psi_{i_s}^s$ that are strings of the form $Q\alpha$ with $Q \in \mathbf{Q} \cup \{\forall, \exists\}$, where α does not contain quantifiers and has variables among v_1, \dots, v_s .

Given an s -tuple $\bar{a} = a_1 \dots a_s$ of elements of a structure A , the *equality type* of \bar{a} is, by definition, the formula

$$\bigwedge \{v_i = v_j \mid a_i = a_j, 1 \leq i < j \leq s\} \wedge \bigwedge \{\neg v_i = v_j \mid a_i \neq a_j, 1 \leq i < j \leq s\}.$$

For any s -tuples $\bar{a}, \bar{b} \in A$ of the same equality type there is a permutation (i.e., an automorphism) of A that maps \bar{a} onto \bar{b} . Hence, for $n \geq 0$, every $L_{\infty\omega}^{\omega}(\mathbf{Q})$ -formula with variables among v_1, \dots, v_s is equivalent in $A_n := \{0, 1, \dots, n\}$ to a disjunction of equality types (the disjunction is finite as there are only finitely many equality types of s -tuples). In particular, there are disjunctions $\chi_i^{(n,s)}$ of equality types of s -tuples such that

$$(1) \quad A_n \models \bigwedge_{1 \leq i \leq l_s} \forall v_1 \dots \forall v_s (\psi_i^s \leftrightarrow \chi_i^{(n,s)}),$$

that is, we have a “local quantifier elimination”. Now, a simple induction on formulas shows:

- (2) For all $\varphi \in L_{\infty\omega}^\omega(\mathbf{Q})$ with variables among v_1, \dots, v_s there is a quantifier-free first-order formula $\chi(v_1, \dots, v_s)$ (and hence, a boolean combination of equality types of s -tuples) such that

$$\bigwedge_{1 \leq i \leq l_s} \forall v_1 \dots \forall v_s (\psi_i^s \leftrightarrow \chi_i^{(n,s)}) \models \forall v_1 \dots \forall v_s (\varphi \leftrightarrow \chi).$$

We set

$$\delta^{(n,s)} := \begin{cases} \text{“there are exactly } n \text{ elements”}, & \text{if } n < s \\ \text{“there are at least } s \text{ elements”}, & \text{if } n \geq s \end{cases}$$

and

$$\varphi^{(n,s)} := \delta^{(n,s)} \wedge \bigwedge_{1 \leq i \leq l_s} \forall v_1 \dots \forall v_s (\psi_i^s \leftrightarrow \chi_i^{(n,s)}).$$

Clearly, $A_n \models \varphi^{(n,s)}$ and $\{\varphi^{(n,s)} \mid n \geq 1\}$ is finite. Now we can show:

- (3) If $B \models \varphi^{(n,s)}$ then for all sentences φ of $L_{\infty\omega}^\omega(\mathbf{Q})$ with variables among v_1, \dots, v_s ,

$$B \models \varphi \quad \text{iff} \quad A_n \models \varphi.$$

In fact, if $\|B\| < s$ then, by $B \models \delta^{(n,s)}$, we have $n = \|B\|$ and hence, $B \cong A_n$; so (3) is trivially true. If $\|B\| \geq s$ then $n \geq s$. For a sentence φ as in (3) choose a boolean combination $\chi(v_1, \dots, v_s)$ of equality types of s -tuples according to (2). By (1), $A_n \models \bigwedge_{1 \leq i \leq l_s} \forall v_1 \dots \forall v_s (\psi_i^s \leftrightarrow \chi_i^{(n,s)})$ and by $B \models \varphi^{(n,s)}$, also $B \models \bigwedge_{1 \leq i \leq l_s} \forall v_1 \dots \forall v_s (\psi_i^s \leftrightarrow \chi_i^{(n,s)})$. Choose $\bar{a} \in A_n^s$ and $\bar{b} \in B^s$ of the same equality type (recall $\|A_n\|, \|B\| \geq s$). Then

$$\begin{aligned} B \models \varphi & \quad \text{iff} \quad B \models \chi[\bar{b}] \\ & \quad \text{iff} \quad A_n \models \chi[\bar{a}] \\ & \quad \text{iff} \quad A_n \models \varphi. \end{aligned}$$

Since $A_n \models \varphi^{(n,s)}$ and every B is isomorphic to some A_n , we see, using (3), that

$$\models_{\text{fin}} \varphi \leftrightarrow \bigvee \{\varphi^{(n,s)} \mid n \geq 1, A_n \models \varphi\}$$

holds for any sentence φ of $L_{\infty\omega}^\omega(\mathbf{Q})$ with variables among v_1, \dots, v_s . As the disjunction on the right side is finite, φ is equivalent to an FO(\mathbf{Q})-sentence.

Altogether, every $L_{\infty\omega}^\omega(\mathbf{Q})$ -sentence is equivalent to an FO(\mathbf{Q})-sentence ψ of the form

$$(4) \quad \psi := \bigvee_{j \in J} (\delta_j \wedge \bigwedge_{i \in I_j} \forall v_1 \dots \forall v_s (\psi_i \leftrightarrow \chi_i)),$$

where $s \geq 1$, the χ_i are quantifier-free, and the δ_j and the ψ_i are of the form $\delta^{(n,s)}$ and ψ_i^s , respectively. Let Ψ be the set of these formulas.

Using a diagonalization procedure, we show that there is a class K of sets such that $K_{<}$ is in PTIME, but K is not axiomatizable by a sentence in Ψ .

For a sentence ψ as in (4) we set

$$\|\psi\| := \text{length of } \psi, \quad s_\psi := s.$$

It is not difficult to define a uniform procedure that, given n and ψ , checks $A_n \models \psi$ in time polynomial in n , $\|\psi\|$, and 2^{s_ψ} . The main point here is that in order to check the validity of $A_n \models \forall v_1 \dots \forall v_{s_\psi} (\psi_i \leftrightarrow \chi_i)$, it suffices to test the validity of $A_n \models (\psi_i \leftrightarrow \chi_i)[\bar{a}]$ for exactly one s_ψ -tuple \bar{a} of each equality type in A_n . This gives a corresponding procedure, because the number of equality types is $O(2^{s_\psi})$ and the (finitely many) quantifiers in $\mathbf{Q} \cup \{\forall, \exists\}$ belong to classes K with $K_{<}$ in PTIME.

We now consider a function $f : \mathbb{N} \rightarrow \Psi$ computable in time polynomial in n that enumerates Ψ in such a way that

$$\|f(n)\| = O(n) \quad \text{and} \quad s_{f(n)} = O(\log n)$$

(it is not hard to see that such an enumeration exists). We then define K by

$$(5) \quad B \in K \quad \text{iff} \quad \text{for some } n, \quad \|B\| = n \text{ and } A_n \not\models f(n).$$

By the preceding remarks, $K_{<}$ is in PTIME; however, by (5), K is not axiomatized by an $f(n)$, that is, K is not axiomatizable by any sentence in Ψ and hence, not by an $L_{\infty\omega}^\omega(\mathbf{Q})$ -sentence.

Now, let τ be an arbitrary vocabulary and let C_τ be the class of structures \mathcal{A} such that $R^{\mathcal{A}}\bar{a}$ for all $R \in \tau$ and $\bar{a} \in A$. As in structures of C_τ any atomic formula $R\bar{t}$ is equivalent to $t_1 = t_1$, the proof above for $\tau = \emptyset$ works for this case too, yielding a class $K \subseteq C_\tau$ such that $K_{<}$ is in PTIME, but K is not axiomatizable in $L_{\infty\omega}^\omega(\mathbf{Q})$. \square

We now prove Theorem 12.2.5. For this purpose we show that for all $s \geq 3$ there are graphs \mathcal{G}_s and $\tilde{\mathcal{G}}_s$ such that the following holds:

- (1) $\mathcal{G}_s \equiv_{L_{\infty\omega}^s(\mathbf{C})} \tilde{\mathcal{G}}_s$.
- (2) There is a class K of graphs such that $K_{<} \in \text{PTIME}$ and K contains all \mathcal{G}_s but no $\tilde{\mathcal{G}}_s$.

This yields the theorem, because, by (1), no $L_{\infty\omega}^\omega(\mathbf{C})$ -sentence φ can axiomatize a class K as in (2).

In section 3.4 we have denoted $L_{\infty\omega}^s(\mathbf{C})$ by $C_{\infty\omega}^s$ and shown that (1) is equivalent to

$$(1') \quad \text{the duplicator wins } C\text{-G}_{\infty}^s(\mathcal{G}_s, \tilde{\mathcal{G}}_s)$$

where $C\text{-G}_{\infty}^s(-, -)$ is the Ehrenfeucht game corresponding to $C_{\infty\omega}^s$.

We define the graphs \mathcal{G}_s and $\tilde{\mathcal{G}}_s$. Let $k := s + 2$. Then $k \geq 5$. Take sets

$$A := \{a_1, \dots, a_k\}, \quad B := \{b_1, \dots, b_k\}, \quad C := \{c_1, \dots, c_k\}$$

of distinct elements, let

$$D := \{S \subseteq \{1, \dots, k\} \mid \|S\| \text{ is even}\},$$

and define the graph $\mathcal{H} := (H, E)$ by

$$\begin{aligned} H &:= A \cup B \cup C \cup D \\ E &:= \{(a_i, c_i), (c_i, a_i), (b_i, c_i), (c_i, b_i) \mid 1 \leq i \leq k\} \\ &\cup \{(a_i, S), (S, a_i) \mid S \in D, i \in S\} \\ &\cup \{(b_i, S), (S, b_i) \mid S \in D, i \notin S\} \\ &\cup \{(c_i, c_{i+1}), (c_{i+1}, c_i) \mid 1 \leq i < k\} \cup \{(c_k, c_1), (c_1, c_k)\}. \end{aligned}$$

In \mathcal{H} , the a_i and b_i have degree $2^{k-2} + 1$, being connected with c_i and exactly 2^{k-2} points in D ; similarly, the $S \in D$ have degree k and the c_i degree 4. (As $k \geq 5$, these degrees are different.) Therefore, any automorphism π of \mathcal{H} maps the sets $A \cup B$, C , and D onto themselves, and if π is the identity on C then $\pi(a_i), \pi(b_i) \in \{a_i, b_i\}$.

Lemma 12.2.7 *Let π be a permutation of $A \cup B \cup C$ which is the identity on C . Then π can be extended to an automorphism of \mathcal{H} iff π interchanges a_i and b_i for an even number of i 's, leaving the other elements in $A \cup B$ fixed.*

Proof. If π interchanges a_i and b_i for $i \in S_0$, then for any automorphism $\tilde{\pi}$ of \mathcal{H} that extends π , and for any $S \in D$ we have

$$(3) \quad \tilde{\pi}(S) = (S \setminus S_0) \cup (S_0 \setminus S)$$

which is in D just in case S_0 is even. If S_0 is even, (3) gives an extension of π which is an automorphism of \mathcal{H} . □

Let V be a set of $s+3 (= k+1)$ elements. We define $\mathcal{H}(V) = (H(V), E^{H(V)})$ as follows: We replace each $v \in V$ by a copy $\mathcal{H}(v)$ of \mathcal{H} , where $a_i(v), b_j(v), \dots$ correspond to a_i, b_j, \dots . Moreover, if $v, w \in V$ and $v \neq w$, we add edges

$$(*) \quad \begin{aligned} &(a(v, w), a(w, v)), (a(w, v), a(v, w)), \\ &(b(v, w), b(w, v)), (b(w, v), b(v, w)) \end{aligned}$$

where

$$\begin{aligned} &\text{for some } i, \quad a(v, w) = a_i(v) \text{ and } b(v, w) = b_i(v), \\ &\text{for some } j, \quad a(w, v) = a_j(w) \text{ and } b(w, v) = b_j(w). \end{aligned}$$

We do this in such a way that for every v and i there are uniquely determined w and j such that $a_i(v)$ is linked to $a_j(w)$ (and hence, $b_i(v)$ is linked to $b_j(w)$).

For $v_0, w_0 \in V$, $v_0 \neq w_0$, we define the “twisted” version $\mathcal{H}(V)^{v_0, w_0}$ just as $\mathcal{H}(V)$, but replacing $(*)$ – only for the points v_0, w_0 – by

$$(a(v_0, w_0), b(w_0, v_0)), (b(w_0, v_0), a(v_0, w_0)), \\ (b(v_0, w_0), a(w_0, v_0)), (a(w_0, v_0), b(v_0, w_0)).$$

Fix $v_0, w_0 \in V, v_0 \neq w_0$. Set

$$\mathcal{G}_s := \mathcal{H}(V), \quad \tilde{\mathcal{G}}_s := \mathcal{H}(V)^{v_0, w_0}.$$

We show that (1') and (2) are satisfied. First we prove that the duplicator wins the game $C\text{-}G_{\infty}^s(\mathcal{G}_s, \tilde{\mathcal{G}}_s)$. Obviously, it is enough to show that he can play in such a way that, after his moves, the following holds:

- (+) There is a $v \in V$ such that $\mathcal{H}(v)$ (in \mathcal{G}_s) is not pebbled and there is an isomorphism $\pi : \mathcal{G}_s^{v, w_0} \cong \tilde{\mathcal{G}}_s$ which respects the pebbles and maps elements of $H(u)$ to elements of $H(u)$ for all $u \in V$.

In the beginning we take $v := v_0$ and π the identity on $H(V)$. Now assume that, before a move of the spoiler, we have v and π according to (+). Let v' be a further point of V different from w_0 such that $\mathcal{H}(v')$ (in \mathcal{G}_s) is not pebbled. Using 12.2.7, choose an automorphism ρ of $\mathcal{H}(v)$ that interchanges $a(v, v')$ and $b(v, v')$ and interchanges $a(v, w_0)$ and $b(v, w_0)$; similarly, choose an automorphism ρ' of $\mathcal{H}(v')$ interchanging $a(v', v)$ and $b(v', v)$, and $a(v', w_0)$ and $b(v', w_0)$. Let σ be the permutation of $H(V)$ that extends $\rho \cup \rho'$ and is the identity outside $H(v) \cup H(v')$. Then it is easy to check that

$$\sigma : \mathcal{G}_s^{v', w_0} \cong \mathcal{G}_s^{v, w_0},$$

i.e., σ shifts the twist from $\{v', w_0\}$ to $\{v, w_0\}$.

Now, if the spoiler chooses a subset X , say, of \mathcal{G}_s , then the duplicator answers by the subset Y of $\tilde{\mathcal{G}}_s$ consisting of the π -images of the elements in $X \setminus H(v)$ and of the $\pi \circ \sigma$ -images of the elements in $X \cap H(v)$. If then the spoiler pebbles an element of Y , the duplicator answers with the corresponding preimage. If $\mathcal{H}(v)$ remains unpebbled, (+) is again satisfied by v and π ; otherwise, it is satisfied by v' and $\pi \circ \sigma : \mathcal{G}_s^{v', w_0} \cong \tilde{\mathcal{G}}_s$.

Before we prove (2), we provide some more information about the graphs \mathcal{G}_s and $\tilde{\mathcal{G}}_s$. $\tilde{\mathcal{G}}_s$ arises from \mathcal{G}_s by a twist corresponding to the pair $\{v_0, w_0\}$. Similarly, we can define twisted versions of \mathcal{G}_s that have twists corresponding to several pairs. We call \mathcal{G} an m -twisted version of \mathcal{G}_s , if there are exactly m twists of this kind. So \mathcal{G}_s is a 0-twisted version of \mathcal{G}_s and $\tilde{\mathcal{G}}_s$ a 1-twisted one.

Claim. Let \mathcal{G} and \mathcal{G}' be m -twisted and m' -twisted versions of \mathcal{G}_s , respectively. Then

$$\mathcal{G} \cong \mathcal{G}' \text{ iff } m - m' \text{ is even.}$$

Proof (of Claim). Let $\pi : \mathcal{G} \cong \mathcal{G}'$. Then, given $v \in V$, π maps $\mathcal{H}(v)$ isomorphically onto some copy $\mathcal{H}(w)$. By (a simple generalization of) 12.2.7, the number of i 's such that $\pi(a_i(v)) \in \{b_1(w), \dots, b_k(w)\}$ is even. This yields

that $m - m'$ is even. Now, let $m - m'$ be even. In the proof of (+) we saw how one can shift twists by suitable isomorphisms, say, a twist corresponding to a pair $\{v', w\}$ to a twist corresponding to a pair $\{v, w\}$. Hence, we may assume that the twists of \mathcal{G} and \mathcal{G}' correspond to pairs that have w_0 as one point. (Note that it may be necessary to shift a twist upon a twist, thus replacing m by $m - 2$ (or, m' by $m' - 2$.) Then by 12.2.7, as $m - m'$ is even, we can find an automorphism of $H(w_0)$ that can be extended to an isomorphism of \mathcal{G} onto \mathcal{G}' .

We now come to a proof of (2). For this purpose we describe a deterministic polynomially time-bounded Turing machine M that accepts all \mathcal{G}_s for $s \geq 3$, but no $\tilde{\mathcal{G}}_s$.

Started with (the ordered version of) a graph \mathcal{G} , M first checks whether there is a number $s \geq 3$ such that \mathcal{G} contains exactly $s + 3$ cycles C_1, \dots, C_{s+3} of length $s + 2$, consisting of points of degree 4 (so-called *c-points*), and whether \mathcal{G} contains no other points of degree 4. If so, M checks

- (i) whether each *c-point* has exactly two neighbours that are no *c-points* (so-called *ab-points*; different *ab-points* with the same *c-neighbour* are called *related*);
- (ii) for each pair $\{a, b\}$ of related *ab-points*, whether a and b are neighbours of exactly one *ab-point* a' and b' , respectively, a' and b' being related via a *c-point* that belongs to a circle different from that of the *c-point* that relates a to b ;
- (iii) for each circle C_i , whether there are exactly 2^{s+1} points (so-called *d-points*) that have exactly $s + 2$ neighbours, these neighbours being *ab-points* that are related to a *c-point* in C_i (we let H_i be the set of these *d-points*, the neighbouring *ab-points*, and the elements of C_i);
- (iv) whether there are no other points or edges.

Finally, M checks for each H_i whether the subgraph \mathcal{H}_i of \mathcal{G} with domain H_i is isomorphic to \mathcal{H} , in the positive case also producing a partition of the *ab-points* into *a-points* and *b-points*. (This can be done in a time polynomial in 2^s and hence, polynomial in $\|\mathcal{G}\|$). If all checks have a positive outcome, M counts the number of twists of \mathcal{G} . If this number is even, M accepts, otherwise it rejects. By the claim above, M accepts the \mathcal{G}_s for $s \geq 3$, but rejects the $\tilde{\mathcal{G}}_s$. \square

12.3 Logical Reductions

In complexity theory there are well-established notions of reducibility, among them the notion of many-one reducibility: Given alphabets \mathbb{A}_i^+ , languages $L_i \subseteq \mathbb{A}_i^+$ ($i = 1, 2$), and a deterministic complexity class \mathcal{C} , the language L_1 is *C-reducible* (or, *C-many-one reducible*) to L_2 if there is a function $f : \mathbb{A}_1^* \rightarrow \mathbb{A}_2^*$ in \mathcal{C} such that

$$w \in L_1 \text{ iff } f(w) \in L_2$$

for all $w \in \mathbb{A}_1^*$. Given a further complexity class C' , a language L is C' -hard with respect to C -reductions if any language in C' is C -reducible to L , and L is C' -complete with respect to C -reductions if it is C' -hard and belongs to C' .

We aim at a logical analogue. For this purpose we need the notion of an interpretation as introduced in section 11.2. For convenience, we shortly repeat the definitions and notations.

Fix a logic \mathcal{L} and vocabularies σ and τ where $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i . An \mathcal{L} -interpretation of σ in τ of width k is given by an $(s+1)$ -tuple of $\mathcal{L}[\tau]$ -formulas

$$\Pi := (\pi_{\text{uni}}(\bar{x}), \pi_{R_1}(\bar{x}_1, \dots, \bar{x}_{r_1}), \dots, \pi_{R_s}(\bar{x}_1, \dots, \bar{x}_{r_s})),$$

where all tuples $\bar{x}, \bar{x}_1, \dots$ are of length k .

For any τ -structure \mathcal{A} , Π induces a σ -structure \mathcal{A}^Π over k -tuples of \mathcal{A} given by

$$\mathcal{A}^\Pi = (\pi_{\text{uni}}^{\mathcal{A}}(-), \pi_{R_1}^{\mathcal{A}}(-, \dots, -), \dots, \pi_{R_s}^{\mathcal{A}}(-, \dots, -)),$$

provided $\pi_{\text{uni}}^{\mathcal{A}}(-)$ is nonempty.² So any \mathcal{L} -interpretation Π of σ in τ induces a map

$$\Pi : \text{Str}[\tau] \rightarrow \text{Str}[\sigma].$$

Exercise 12.3.1 Let Π_1 be an \mathcal{L} -interpretation of σ in τ and Π_2 be an \mathcal{L} -interpretation of τ in ρ . Show that there is an \mathcal{L} -interpretation Π of σ in ρ such that for all ρ -structures \mathcal{A} ,

$$\mathcal{A}^\Pi = (\mathcal{A}^{\Pi_2})^{\Pi_1}. \quad \square$$

Taking the maps Π as candidates for the maps linked to logical reducibility, we are led to the following analogue of the notion of reducibility in complexity theory.

Definition 12.3.2 Let $K \subseteq \text{Str}[\tau]$ and $L \subseteq \text{Str}[\sigma]$. K is \mathcal{L} -reducible to L , for short: $K \leq_{\mathcal{L}} L$, if there is an \mathcal{L} -interpretation Π of σ in τ such that for all τ -structures \mathcal{A} ,

$$\mathcal{A} \in K \text{ iff } \mathcal{A}^\Pi \in L.$$

We write $K \leq_{\mathcal{L}}^k L$ to indicate that $K \leq_{\mathcal{L}} L$ can be witnessed by an interpretation of width k . □

Exercise 12.3.3 Show that $\leq_{\mathcal{L}}$ is transitive. □

² As in section 11.2, we tacitly ignore the case $\pi_{\text{uni}}^{\mathcal{A}}(-) = \emptyset$ as it will not be of importance; moreover, note that $\pi_{R_i}^{\mathcal{A}}(-, \dots, -)$ stands for $\pi_{R_i}^{\mathcal{A}}(-, \dots, -) \cap (\pi_{\text{uni}}^{\mathcal{A}}(-))^{r_i}$.

Exercise 12.3.4 Let $K \subseteq \text{Str}[\tau]$ be a class of ordered structures and $L \subseteq \text{Str}[\sigma]$. Show that $K \leq_{\mathcal{L}} L$ implies $K \leq_{\mathcal{L}} L_{<}$. \square

An important example of first-order reductions is given in the following consequence of 11.2.5. In fact, taking as L the class K^{II} (II as in 11.2.5) we have

Proposition 12.3.5 *Let τ be an arbitrary vocabulary. Then, for all $K \subseteq \text{Str}[\tau]$ there is a class $L \subseteq \text{GRAPH}$ such that $K \leq_{\text{FO}} L$. Moreover, L can be chosen such that $K_{<} \in \text{PTIME}$ iff $L_{<} \in \text{PTIME}$.* \square

We study the relationship between reductions and Lindström quantifiers. First, there is a difference: Interpretations II contain a formula $\pi_{\text{uni}}(\bar{x})$ giving the new universe, whereas Lindström quantifiers do not allow to “relativize”, they always refer to the whole universe of the structure in question. To overcome this difference, we proceed as follows:

Let $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i and let U be a new unary relation symbol. For $L \subseteq \text{Str}[\sigma]$ we introduce the class $\text{rel-}L \subseteq \text{Str}[\sigma \cup \{U\}]$ by

$$\text{rel-}L := \{(A, U^A, R_1^A, \dots, R_s^A) \mid (U^A, R_1^A, \dots, R_s^A) \in L\};$$

more precisely, $\text{rel-}L$ is the class

$$\{(A, U^A, R_1^A, \dots, R_s^A) \mid (U^A, R_1^A \cap (U^A)^{r_1}, \dots, R_s^A \cap (U^A)^{r_s}) \in L\}.$$

For $k \geq 1$ set $\sigma(k) := \{R_1(k), \dots, R_s(k)\}$ with $r_i \cdot k$ -ary $R_i(k)$ and let $U(k)$ be an k -ary relation symbol. Identifying $r \cdot k$ -tuples over A with r -tuples over A^k as in section 12.1, we “relativize” the definition of vectorization (cf. 12.1.6) as follows:

Definition 12.3.6 Let $L \subseteq \text{Str}[\sigma]$.

(a) $\text{rel-}k\text{-}L$, the *relativized k -vectorization* of L , is the following class of $\{U(k)\} \cup \sigma(k)$ -structures:

$$\begin{aligned} \text{rel-}k\text{-}L &:= \{(A, W, S_1, \dots, S_s) \mid (A^k, W, S_1, \dots, S_s) \in \text{rel-}L\} \\ & (= \{(A, W, S_1, \dots, S_s) \mid W \subseteq A^k \text{ and } (W, S_1, \dots, S_s) \in L\}). \end{aligned}$$

Note that $\text{rel-}L$ is $\text{rel-}1\text{-}L$.

(b) The Lindström quantifier $Q_{\text{rel-}k\text{-}L}$ is called the *relativized k -vectorization* of the quantifier Q_L . We set $\text{rel-}\omega\text{-}\mathbf{Q}_L := \{Q_{\text{rel-}k\text{-}L} \mid k \geq 1\}$. \square

Example 12.3.7 For the class CONN of connected graphs we have

$$A \models Q_{\text{rel-}k\text{-}\text{CONN}} \bar{x}, \bar{y} [\varphi(\bar{x}), \psi(\bar{x}, \bar{y})] \text{ iff } (\varphi^A(-), \psi^A(-, -)) \in \text{CONN}. \quad \square$$

As a simple observation we have (compare 12.1.6 for the definition of L^k):

Proposition 12.3.8 *Let L be a class of structures and \mathcal{L} a logic. Then for all $k \geq 1$,*

$$\mathcal{L}(Q_{L^k}) \leq \mathcal{L}(Q_{\text{rel-}k\text{-}L}).$$

Proof. If, for instance, $k = 2$ and L is a class of $\{R\}$ -structures with r -ary R , the formulas

$$Q_{L^2} x_1 x'_1 \dots x_r x'_r \psi(x_1 x'_1, \dots, x_r x'_r),$$

$$Q_{\text{rel-}2\text{-}L} x x' x_1 x'_1 \dots x_r x'_r [x = x \wedge x' = x', \psi(x_1 x'_1, \dots, x_r x'_r)]$$

are equivalent. \square

The direction “ \geq ” may be wrong:

Exercise 12.3.9 Let P be unary and $L := \{(A, P^A) \mid \|P^A\| \leq \frac{1}{2}\|A\|\}$. Show that

$$A \models Q_{\text{rel-}L} x, y [\varphi(x), \psi(y)] \quad \text{iff} \quad \varphi^A(-) \neq \emptyset \text{ and} \\ \|\psi^A(-) \cap \varphi^A(-)\| \leq \frac{1}{2}\|\varphi^A(-)\|$$

and that $\text{FO}(Q_L) < \text{FO}(Q_{\text{rel-}L})$. \square

Exercise 12.3.10 Let I and P be the classes introduced in 12.1.9. Show that $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_I) \equiv \text{FO}(\mathbf{Q}_I^{\omega}) \equiv \text{FO}(\text{IFP})$ and $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_P) \equiv \text{FO}(\mathbf{Q}_P^{\omega}) \equiv \text{FO}(\text{PFP})$.

Exercise 12.3.11 Prove $\mathcal{L}(\text{rel-}\omega\text{-}\mathbf{Q}_L) \equiv \mathcal{L}(\text{rel-}\omega\text{-}\mathbf{Q}_{L_{<}})$ on ordered structures (compare 12.1.2(f)). \square

Exercise 12.3.12 Let K be a class such that $K_{<} \in \text{PTIME}$. Show that $\text{Mod}(\psi)_{<} \in \text{PTIME}$ for any sentence $\psi \in \text{FO}(\text{IFP})(\text{rel-}\omega\text{-}\mathbf{Q}_K)$. Moreover show: If $\text{FO}(\text{IFP})(\text{rel-}\omega\text{-}\mathbf{Q}_K)$ or $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_K)$ strongly captures PTIME , then it effectively strongly captures PTIME . \square

For \mathcal{L} a logic and \mathbf{Q} a set of Lindström quantifiers, we denote by $\mathcal{L}[\mathbf{Q}]$ the set of formulas of $\mathcal{L}(\mathbf{Q})$ that are of the form

$$Q \bar{x}_1, \dots, \bar{x}_s [\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)]$$

where $Q \in \mathbf{Q}$ and ψ_1, \dots, ψ_s are \mathcal{L} -formulas (thus the ψ_i do not contain the new quantifiers). Then a simple, but crucial link between reducibility and quantifiers can be stated as follows:

Proposition 12.3.13 *Let \mathcal{L} be a logic, $K \subseteq \text{Str}[\tau]$ and $L \subseteq \text{Str}[\sigma]$. Then for $k \geq 1$ the following are equivalent:*

- (i) $K \leq_{\mathcal{L}}^k L$.
- (ii) *There is a sentence φ in $\mathcal{L}[Q_{\text{rel-}k\text{-}L}]$ of vocabulary τ such that $K = \text{Mod}(\varphi)$.*

Proof. Assume $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i . To prove (i) \Rightarrow (ii), suppose that $K \leq_{\mathcal{L}}^k L$ via

$$\Pi := (\pi_{\text{uni}}(\bar{x}), \pi_{R_1}(\bar{x}_1, \dots, \bar{x}_{r_1}), \dots, \pi_{R_s}(\bar{x}_1, \dots, \bar{x}_{r_s})).$$

Then for any τ -structure \mathcal{A} we have:

$$\begin{aligned} \mathcal{A} \in K & \text{ iff } \mathcal{A}^\Pi = (\pi_{\text{uni}}^{\mathcal{A}}(-), \pi_{R_1}^{\mathcal{A}}(-, \dots, -), \dots, \pi_{R_s}^{\mathcal{A}}(-, \dots, -)) \in L \\ & \text{ iff } (A, \pi_{\text{uni}}^{\mathcal{A}}(-), \pi_{R_1}^{\mathcal{A}}(-, \dots, -), \dots, \pi_{R_s}^{\mathcal{A}}(-, \dots, -)) \in \text{rel-}k\text{-}L \\ & \text{ iff } \mathcal{A} \models Q_{\text{rel-}k\text{-}L} \bar{x}, \bar{x}_1, \dots, \bar{x}_{r_s} [\pi_{\text{uni}}, \pi_{R_1}, \dots, \pi_{R_s}]. \end{aligned}$$

Concerning the direction from (ii) to (i), assume that K is axiomatized by the sentence

$$Q_{\text{rel-}k\text{-}L} \bar{x}, \bar{x}_1, \dots, \bar{x}_{r_s} [\pi_{\text{uni}}, \pi_{R_1}, \dots, \pi_{R_s}]$$

with \mathcal{L} -formulas $\pi_{\text{uni}}, \pi_{R_1}, \dots, \pi_{R_s}$. Then a rearrangement of the equivalences above shows that $K \leq_{\mathcal{L}}^k L$ via $\Pi = (\pi_{\text{uni}}, \pi_{R_1}, \dots, \pi_{R_s})$. \square

We now transfer the notions of hardness and completeness from complexity theory to logic.³

Definition 12.3.14 Let \mathcal{C} be a complexity class, $L \subseteq \text{Str}[\sigma]$, and \mathcal{L} a logic.

- (a) L is \mathcal{C} -hard with respect to \mathcal{L} -reductions if $K \leq_{\mathcal{L}} L$ for all $K \in \mathcal{C}$.
- (b) L is \mathcal{C} -complete with respect to \mathcal{L} -reductions if L is \mathcal{C} -hard with respect to \mathcal{L} -reductions and $L_{<} \in \mathcal{C}$. \square

Corollary 12.3.15 Let \mathcal{C}, L , and \mathcal{L} be as in the preceding definition.

- (a) The following are equivalent:
 - (i) L is \mathcal{C} -hard with respect to \mathcal{L} -reductions.
 - (ii) $\mathcal{C} \subseteq \{\text{Mod}(\varphi) \mid \varphi \in \mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_L]\}$.
- (b) If \mathcal{C} is closed under \mathcal{L} -reducibility (i.e., for classes K_1, K_2 of ordered structures, $K_1 \leq_{\mathcal{L}} K_2$ and $K_2 \in \mathcal{C}$ imply $K_1 \in \mathcal{C}$), then the following are equivalent:
 - (i) L is \mathcal{C} -complete with respect to \mathcal{L} -reductions.
 - (ii) $\mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_L]$ captures \mathcal{C} .

Proof. (a) is immediate from the preceding proposition. For the implication (ii) \Rightarrow (i) in (b) note that $L_{<}$ (as L) is $\text{FO}[Q_{\text{rel-}L}]$ -axiomatizable and therefore $\mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_L]$ -axiomatizable, hence $L_{<} \in \mathcal{C}$ by (ii). The rest is clear by (a). For (i) \Rightarrow (ii) assume that the class K of ordered structures is axiomatizable in $\mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_L]$. Then $K \leq_{\mathcal{L}} L$ by 12.3.13 and hence, $K \leq_{\mathcal{L}} L_{<}$ (cf. 12.3.4). Since $L_{<} \in \mathcal{C}$ (by (i)) and \mathcal{C} is closed under \mathcal{L} -reducibility, we get $K \in \mathcal{C}$. The rest is clear by (a). \square

Corollary 12.3.16 Let \mathbf{I} and \mathbf{P} be the classes introduced in the proof of 12.1.9. Then \mathbf{I} is PTIME -complete with respect to FO -reductions and \mathbf{P} is PSPACE -complete with respect to FO -reductions.

³ Recall our convention that only classes of ordered structures are considered as members of complexity classes.

Proof. By the proof of 12.1.9, $\text{FO}(\text{IFP}) \leq \text{FO}[\mathbf{Q}_I^\varphi]$. Therefore, $\text{FO}(\text{IFP}) \leq \text{FO}[\mathbf{Q}_I^\omega] \leq \text{FO}[\text{rel-}\omega\text{-}\mathbf{Q}_I] \leq \text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_I) \leq \text{FO}(\text{IFP})$ (by 12.3.10). Since $\text{FO}(\text{IFP})$ captures PTIME, so does $\text{FO}[\text{rel-}\omega\text{-}\mathbf{Q}_I]$. Thus, I is PTIME-complete with respect to FO-reductions by part (b) of the preceding corollary. The proof for P is similar. \square

We are now in a position to show the result announced in section 11.1 that in case PTIME is effectively strongly captured by any logic at all, it can already be captured by a very familiar one, namely by an extension of first-order logic by means of a single vectorized Lindström quantifier.

A class L is *strongly PTIME-complete with respect to FO-reductions* if $L_{<}$ is in PTIME and $K \leq_{\text{FO}} L$ for any class K of structures with $K_{<}$ in PTIME (in Definition 12.3.14 we only considered classes K of ordered structures). The main result is:

Theorem 12.3.17 *The following are equivalent:*

- (i) *There is a logic that effectively strongly captures PTIME.*
- (ii) *There is a class L of structures such that $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_L)$ effectively strongly captures PTIME.*
- (iii) *There is a class L of structures which is strongly PTIME-complete with respect to FO-reductions.*

Moreover, a class L satisfies (ii) just in case it satisfies (iii).

The theorem gets false if “strongly PTIME-complete” is replaced by “PTIME-complete”: With respect to FO-reductions the class I is PTIME-complete (cf. 12.3.16) but not strongly PTIME-complete (otherwise, by the last statement of the theorem, $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_I) \equiv \text{FO}(\text{IFP})$ (cf. the proof of 12.3.16) would effectively strongly capture PTIME).

Proof (of 12.3.17). We show (iii) \Rightarrow (ii) and (i) \Rightarrow (iii), the implication (ii) \Rightarrow (i) being trivial.

(iii) \Rightarrow (ii): Let the class L be given by (iii) and let K be any class of structures such that $K_{<} \in \text{PTIME}$. By (iii), $K \leq_{\text{FO}} L$ and hence, by 12.3.13, K is axiomatizable in $\text{FO}[\text{rel-}\omega\text{-}\mathbf{Q}_L] \subseteq \text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_L)$. Moreover, by induction one can effectively assign to every $\varphi \in \text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_L)$ a pair (M, d) , where M is a deterministic x^d time-bounded machine accepting $\text{Mod}(\varphi)_{<}$. Hence, $\text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_L)$ effectively strongly captures PTIME (note that $\text{Mod}(\psi)_{<} \in \text{PTIME}$ for $\psi \in \text{FO}(\text{rel-}\omega\text{-}\mathbf{Q}_L)$, cf. 12.3.12).

(i) \Rightarrow (iii): It suffices to define a class L of structures such that $L_{<}$ is in PTIME and such that for every class K of $\{E\}$ -structures (where E is binary) with $K_{<}$ in PTIME we have $K \leq_{\text{FO}} L$. Then 12.3.5 together with the transitivity of \leq_{FO} (cf. 12.3.3) yields the strong PTIME-completeness of L .

Let \mathcal{L} be a logic that effectively strongly captures PTIME. We may assume that the set $\mathcal{L}\{\{E\}\}_0$ of sentences of \mathcal{L} of vocabulary $\{E\}$ is \mathbb{N} . Furthermore,

let g be a computable function assigning to $i = \varphi \in \mathcal{L}\{\{E\}\}_0$ a pair (M, d) such that M is deterministic and x^d time-bounded and accepts $\text{Mod}(\varphi)_<$. Let M_g be a deterministic machine that calculates g .

We set

$$\sigma := \{V, E, P, Q\}$$

with unary V, Q and binary E, P and define L as the class of σ -structures A such that:

- (1) $E^A \subseteq V^A \times V^A$.
- (2) P^A is a preordering of A (i.e., it is reflexive and transitive) such that the equivalence classes of the equivalence relation given by

$$a \text{ and } b \text{ are equivalent iff } P^A ab \text{ and } P^A ba$$

are linearly ordered with respect to the ordering \preceq^A induced by P^A .

- (3) Q^A is an equivalence class, say, Q^A is the i -th class with respect to \preceq^A .
- (4) The length of the run that M_g performs on the input $1 \dots 1$ of length i to calculate $g(i) = (M, d)$ is $\leq \|A\|$.
- (5) $\|A\| \geq \|V^A\|^d$.
- (6) M accepts all ordered representations of (V^A, E^A) .

$L_<$ is in PTIME. Namely, in order to check in time polynomial in $\|A\|$ whether an ordered $\sigma_<$ -structure $(A, <^A)$ with $A = \{1, \dots, n\}$ belongs to $L_<$, we

- check (1)
- check whether P^A satisfies (2) and whether Q^A is an equivalence class and determine i
- calculate $g(i) = (M, d)$, thereby checking (4)
- check (5)
- check whether M accepts $(V^A, E^A, <^A \cap (V^A \times V^A))$ (which is equivalent to a check of (6), as M is x^d time-bounded and accepts $\text{Mod}(i)_<$).

Now, let K be a class of $\{E\}$ -structures such that $K_<$ is in PTIME. We show that $K \leq_{\text{FO}} L$. For this purpose let $i = \varphi \in \mathcal{L}\{\{E\}\}_0$ be such that $\text{Mod}(\varphi) = K$. Let $g(i) = (M, d)$. Then the deterministic machine M is x^d time-bounded and accepts $K_<$. Choose l such that there are at least $\max\{i, d\}$ many equality types of l elements (cf. section 12.2) and such that $l^l \geq$ the length of the run that M_g performs on the input $1 \dots 1$ of length i to calculate $g(i)$. We only take care of $\{E\}$ -structures of cardinality $\geq l$.

We show $K \leq_{\text{FO}}^l L$ by defining the following FO-interpretation Π of σ in $\{E\}$ of width l (\bar{x} and \bar{y} are of length l):

$$\begin{aligned}
 \pi_{\text{uni}}(\bar{x}) &:= x_1 = x_1 \\
 \pi_V(\bar{x}) &:= x_1 = x_2 = \dots = x_l \\
 \pi_P(\bar{x}, \bar{y}) &:= \text{“the equality type of } \bar{x} \text{ is smaller than that of } \bar{y}\text{”} \\
 &\quad \text{(in a fixed, first-order definable ordering of the finitely} \\
 &\quad \text{many equality types of } l \text{ elements)} \\
 \pi_Q(\bar{x}) &:= \text{“} \bar{x} \text{ satisfies the } i\text{-th equality type”} \\
 \pi_E(\bar{x}, \bar{y}) &:= x_1 = \dots = x_l \wedge y_1 = \dots = y_l \wedge E x_1 y_1.
 \end{aligned}$$

Then for all $\{E\}$ -structures \mathcal{B} of cardinality $\geq l$,

$$\mathcal{B} \in K \quad \text{iff} \quad \mathcal{B}^\Pi \in L. \quad \square$$

To conclude this section, we give some further applications of Corollary 12.3.15 for $\mathcal{C} = \text{NPTIME}$.

Let $< \in \sigma$ and let $L \subseteq \mathcal{O}[\sigma]$ be a class of ordered structures that is NPTIME-complete with respect to logspace reductions, i.e., $L \in \text{NPTIME}$ and every NPTIME class K of ordered structures, say of vocabulary τ , is reducible to L in the following sense: There is a log space-bounded⁴ (τ, σ) -transducer M (cf. Appendix 7.5.1) such that for all τ -structures \mathcal{A} ,

$$\mathcal{A} \in K \quad \text{iff} \quad M(\mathcal{A}) \in L$$

where $M(\mathcal{A})$ is the output of M started with \mathcal{A} . From complexity theory we know that an example of such an L is given by $\text{HAM}_{<}$ where

$$\text{HAM} := \{ \mathcal{G} \mid \mathcal{G} \text{ is a directed graph that has a Hamiltonian circuit} \}.$$

Proposition 12.3.18 *For classes K and L of ordered structures the following are equivalent:*

- (i) K is logspace reducible to L .
- (ii) $K \leq_{\text{FO(DTC)}} L$.

Proof. For the implication (ii) \Rightarrow (i) let Π be an FO(DTC)-interpretation witnessing $K \leq_{\text{FO(DTC)}} L$. Since satisfiability for FO(DTC)-formulas can be checked by log space-bounded machines, there is a (τ, σ) -transducer that, given \mathcal{A} , writes an encoding of \mathcal{A}^Π on the output tapes.

For (i) \Rightarrow (ii) suppose that $L \subseteq \mathcal{O}[\sigma]$ where $\sigma = \{R_1, \dots, R_s\}$ with r_i -ary R_i . Let $K \subseteq \mathcal{O}[\tau]$ and let M be a log space-bounded (τ, σ) -transducer such that for all $\mathcal{A} \in \mathcal{O}[\tau]$,

$$(1) \quad \mathcal{A} \in K \quad \text{iff} \quad M(\mathcal{A}) \in L.$$

By 7.5.23 there are $k \geq 1$ and formulas $\chi_{\text{uni}}(\bar{x})$ and $\chi_{R_i}(\bar{x}_1, \dots, \bar{x}_{r_s})$ of FO(DTC) with $\bar{x}, \bar{x}_1, \dots, \bar{x}_{r_s}$ of length k such that for all $\mathcal{A} \in \mathcal{O}[\tau]$,

⁴ Recall that log space-bounded means $c \cdot \log$ space-bounded for a suitable $c > 0$.

$$(2) \quad M(\mathcal{A}) \cong (\chi_{\text{uni}}^{\mathcal{A}}(-), \chi_{R_1}^{\mathcal{A}}(-, \dots, -), \dots, \chi_{R_s}^{\mathcal{A}}(-, \dots, -)).$$

Let Π be the FO(DTC)-interpretation $(\chi_{\text{uni}}, \chi_{R_1}, \dots, \chi_{R_s})$. By(1) and (2), we have $\mathcal{A} \in K$ iff $\mathcal{A}^\Pi \in L$; hence, $K \leq_{\text{FO(DTC)}}^k L$. \square

Since NPTIME is closed under logspace reductions, we get by 12.3.15(b) and the preceding proposition:

Corollary 12.3.19 *For a class L of ordered structures the following are equivalent:*

- (i) L is NPTIME-complete with respect to logspace reductions.
- (ii) L is NPTIME-complete with respect to FO(DTC)-reductions.
- (iii) FO(DTC)[rel- ω - \mathbf{Q}_L] captures NPTIME. \square

Corollary 12.3.20 $\text{HAM}_{<}$ is NPTIME-complete with respect to FO(DTC)-reductions. \square

Corollary 12.3.21 FO(DTC)[rel- ω - $\mathbf{Q}_{\text{HAM}_{<}}$] captures NPTIME. \square

In some cases one can do better. For example, NPTIME is captured already by FO[rel- ω - $\mathbf{Q}_{\text{HAM}_{<}}$] (in the presence of min and max even quantifier-free formulas suffice); moreover, according to the following exercise we can replace $\text{HAM}_{<}$ by HAM.

Exercise 12.3.22 Show that $\mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_L] \equiv \mathcal{L}[\text{rel-}\omega\text{-}\mathbf{Q}_{L_{<}}]$ on ordered structures (compare 12.1.2(f)). \square

In a typical first-order prefix, for example a prefix as in

$$(1) \quad \forall x \exists y \forall u \exists v \varphi(x, y, u, v),$$

the existentially quantified variables depend on all the preceding universally quantified variables. So (1) means (using function symbols f (unary) and g (binary))

$$(2) \quad \exists f \exists g \forall x \forall u \varphi(x, f(x), u, g(x, u)).$$

Due to the linear character of our notation of formulas it is not possible to arrange the quantifiers in the prefix in (1) in such a way that y only depends on x and v only depends on u . Henkin introduced quantifiers which allow such dependencies. Here we only consider the simplest one, the one we just have described. It is written as $\left(\begin{smallmatrix} \forall x \exists y \\ \forall u \exists v \end{smallmatrix} \right)$, and $\left(\begin{smallmatrix} \forall x \exists y \\ \forall u \exists v \end{smallmatrix} \right) \varphi(x, y, u, v)$ means $\exists f \exists g \forall x \forall u \varphi(x, f(x), u, g(u))$. We turn $\left(\begin{smallmatrix} \forall x \exists y \\ \forall u \exists v \end{smallmatrix} \right)$ into a Lindström quantifier by introducing, for $\sigma = \{R\}$ with a 4-ary relation symbol R , the class

$$\text{HENK} := \{(A, R^A) \mid (A, R^A) \models \exists f \exists g \forall x \forall u R x f(x) u g(u)\}.$$

$\text{FO}(Q_{\text{HENK}})$ is strictly stronger than FO. For example, with $\tau = \{U, V\}$, where U, V are unary, and

$$\begin{aligned} \varphi &:= \left(\begin{array}{l} \forall x \exists y \\ \forall u \exists v \end{array} \right) ((u = y \rightarrow v = x) \wedge (Ux \leftrightarrow Vy)) \\ &= \text{“}\exists f \exists g \forall x \forall u ((u = f(x) \rightarrow g(u) = x) \wedge (Ux \leftrightarrow Vf(x)))\text{”} \\ &= Q_{\text{HENK}}xyuv((u = y \rightarrow v = x) \wedge (Ux \leftrightarrow Vy)) \end{aligned}$$

we have $\text{Mod}(\varphi) = \{(A, U^A, V^A) \mid \|U^A\| \leq \|V^A\|\}$, a class that is not axiomatizable in FO.

Proposition 12.3.23 *For $\mathbf{Q} = \text{rel-}\omega\text{-}\mathbf{Q}_{\text{HENK}}$ and for $\mathbf{Q} = \text{rel-}\omega\text{-}\mathbf{Q}_{\text{HENK}<}$ the logic $\text{FO}(\text{DTC})[\mathbf{Q}]$ captures NPTIME.*

Proof. By 12.3.22 we may restrict ourselves to $\text{HENK}_{<}$. Let 3COL be the class of 3-colourable graphs. (A graph is 3-colourable, if its vertices can be coloured with three colours such that no two adjacent vertices have the same colour.) $\text{3COL}_{<}$ is NPTIME-complete with respect to logspace reductions and hence, with respect to FO(DTC)-reductions (cf. 12.3.19). We show

$$(*) \quad \text{3COL}_{<} \leq_{\text{FO}(\text{DTC})} \text{HENK}_{<}$$

Since $\text{HENK}_{<}$ is in NPTIME, (*) yields that $\text{HENK}_{<}$ is NPTIME-complete with respect to FO(DTC)-reductions. By 12.3.19, $\text{FO}(\text{DTC})[\text{rel-}\omega\text{-}\mathbf{Q}_{\text{HENK}<}]$ captures NPTIME.

We show (*), even $\text{3COL}_{\leq \text{FO}} \text{HENK}$: We use a unary function f to represent three colours: a vertex a has the first colour if

$$f(a) = a,$$

the second colour if

$$f(a) \neq a \text{ and } Eaf(a),$$

and the third colour if

$$f(a) \neq a \text{ and not } Eaf(a).$$

Then an easy argumentation shows that a graph $\mathcal{G} = (G, E^{\mathcal{G}})$ is 3-colourable iff it is of cardinality ≤ 3 or there is a function $f : G \rightarrow G$ such that for all a, b with $E^{\mathcal{G}}ab$ it is not the case that (i), (ii), or (iii) where

- (i) $f(a) = a$ and $f(b) = b$;
- (ii) $f(a) \neq a, Eaf(a)$ and $f(b) \neq b, Ebf(b)$;
- (iii) $f(a) \neq a$, not $Eaf(a)$ and $f(b) \neq b$, not $Ebf(b)$.

Therefore, there is a formula $\varphi_{\text{3COL}}(x, y, u, v) \in \text{FO}[\{E\}]$ such that

$$\text{Mod}(Q_{\text{HENK}}xyuv\varphi_{\text{3COL}}(x, y, u, v)) = \text{3COL};$$

for instance, we can take

$$\begin{aligned} \varphi_{3\text{COL}}(x, y, u, v) &:= \varphi_{\leq 3} \vee ((x = u \rightarrow y = v) \wedge (Exu \rightarrow \\ &\quad \neg((y = x \wedge v = u) \vee (y \neq x \wedge Exy \wedge v \neq u \wedge Euv) \\ &\quad \vee (y \neq x \wedge \neg Exy \wedge v \neq u \wedge \neg Euv))))). \end{aligned}$$

Hence, $3\text{COL} \leq_{\text{FO}} \text{HENK}$ via the interpretation given by

$$\begin{aligned} \pi_{\text{uni}}(x) &:= x = x \\ \pi_R(x, y, u, v) &:= \varphi_{3\text{COL}}(x, y, u, v). \end{aligned}$$

□

12.4 Quantifiers and Oracles

In the preceding section we have encountered complexity classes captured by logics of the form $\mathcal{L}(\text{rel-}\omega\text{-}\mathbf{Q}_K)$. In the following we deal with a question in the opposite direction: Let \mathcal{L} be a logic, K a class of structures, and \mathcal{C} a complexity class. Then:

(+) If \mathcal{L} captures \mathcal{C} , what complexity class is captured by $\mathcal{L}(\text{rel-}\omega\text{-}\mathbf{Q}_K)$?

Interprete a formula φ , say, of $\mathcal{L}(\mathbf{Q}_K)$ as a program to evaluate the meaning of φ , addressing a device which can evaluate instructions of the form

$$Q_K \bar{x}[\varphi_1, \dots, \varphi_s]$$

in one step, once $\varphi_1, \dots, \varphi_s$ have been evaluated. There is a concept of such a device in complexity theory that fits into this frame, Turing machines with *oracles*. We let \mathcal{C}^K denote the complexity class defined by Turing machines with resource-bound according to \mathcal{C} that can use an oracle for K to decide membership for K . Then the natural answer to (+) would be

$$\mathcal{L}(\text{rel-}\omega\text{-}\mathbf{Q}_K) \text{ captures } \mathcal{C}^K.$$

As it has turned out, this answer is true in various cases. However, one has to choose the right oracle machine model in each particular case. We only will present one important group of results to give an impression of what is going on. First, we introduce an oracle model (essentially Turing machines with one oracle tape) that meets our purposes.

Let σ and τ be vocabularies with $< \in \sigma$ and $< \in \tau$. A *Turing machine M for τ -structures with oracle tapes for σ -structures*, for short, a (τ, σ) -*oracle machine*, is a (τ, σ) -Turing machine (cf. Appendix 7.5.1), that is, a deterministic Turing machine with input tapes for τ -structures, with work tapes, and with output tapes for σ -structures (now called *oracle tapes*). In addition to

the instructions of a (τ, σ) -Turing machine it has so-called *oracle instructions*, that is, instructions of the form

$$(1) \quad sb_0 \dots b_{k+l} c_1 \dots c_m \rightarrow s' s''.$$

Given an “oracle class” O of σ -structures, (1) has the meaning

“If you are in a configuration corresponding to $sb_0 \dots b_{k+l} c_1 \dots c_m$, do the following: if the inscription on the oracle tapes encodes a structure in O , go into state s' ; otherwise, go into state s'' . Delete the inscription on the oracle tapes and move the heads to the initial position.”

Configurations which ask for an oracle instruction are called *oracle configurations*. Note that in an oracle configuration the inscriptions on the oracle tapes need not encode a σ -structure.

We write (M, O) to indicate that M works with O as the class fixing the meaning of the oracle instructions.

An oracle instruction is carried out in one step. By definition, space bounds refer only to work tapes; in particular, space used on the oracle tapes does not count. If (M, O) is $s(n)$ space-bounded then – in case (M, O) stops – the number of cells on the oracle tapes that are scanned during the run is $\leq d^{s(n)}$ for a certain constant $d > 0$.

Occasionally we shall also use transducers with oracle tapes, so-called *oracle transducers*; they possess both output tapes and oracle tapes, the latter ones serving the same purpose and working in the same way as with oracle machines. It should not be difficult for the reader to provide an exact definition.

Let \mathcal{C} be a complexity class and \mathcal{K} a class of classes of ordered structures. $\mathcal{C}^{\mathcal{K}}$, the *relativization* of \mathcal{C} to \mathcal{K} , consists of those classes of ordered structures that are accepted by some oracle machine (M, O) with $O \in \mathcal{K}$ that is resource-bounded according to \mathcal{C} . In particular, $\text{LOGSPACE}^{\mathcal{K}}$ consists of those classes of ordered structures that are accepted by some log space-bounded oracle machine (M, O) with $O \in \mathcal{K}$. If L is a class of ordered structures, we write \mathcal{C}^L for $\mathcal{C}^{\{L\}}$.

The following proposition which is stated for the case we are interested in exhibits the relationship between oracles and quantifiers.

Proposition 12.4.1 *Suppose that O is a class of ordered σ -structures and φ an $\text{FO}(\text{DTC})[\text{rel-}\omega\text{-}\mathbf{Q}_O]$ -sentence. Then there is a log space-bounded oracle machine (M, O) that decides $\text{Mod}(\varphi)$.⁵*

Proof. Let

$$\varphi = Q_{\text{rel-}k\text{-}O} \bar{x}, \bar{x}_1, \dots, \bar{x}_{r_s} [\varphi_{\text{uni}}(\bar{x}), \varphi_{R_1}(\bar{x}_1, \dots, \bar{x}_{r_1}), \dots, \varphi_{R_s}(\bar{x}_1, \dots, \bar{x}_{r_s})]$$

⁵ In particular, the logarithmic space bound applies to rejecting runs, too.

with $\varphi_{\text{uni}}, \dots, \varphi_{R_s} \in \text{FO}(\text{DTC})$. By 7.4.1 there are log space-bounded Turing machines that evaluate $\varphi_{\text{uni}}, \dots, \varphi_{R_s}$. The intended machine writes the “structure” induced by $\varphi_{\text{uni}}, \dots, \varphi_{R_s}$ on the oracle tapes and invokes its oracle to decide whether it belongs to O . \square

We prove the following well-known result from complexity theory.

Proposition 12.4.2 *Assume that L is a class of ordered structures that is NPTIME-complete with respect to logspace reductions. Then*

$$\text{LOGSPACE}^L = \text{LOGSPACE}^{\text{NPTIME}}.$$

In particular, we have

$$\text{LOGSPACE}^{\text{HAM}<} = \text{LOGSPACE}^{\text{HENK}<} = \text{LOGSPACE}^{\text{NPTIME}}.$$

Proof. Obviously, we have $\text{LOGSPACE}^L \subseteq \text{LOGSPACE}^{\text{NPTIME}}$. To prove the converse inclusion, let $K \subseteq \mathcal{O}[\tau]$ and $O \subseteq \mathcal{O}[\sigma]$ with $O \in \text{NPTIME}$, and let (M_K, O) be a log space-bounded (τ, σ) -oracle machine that accepts K . We have to show that $K \in \text{LOGSPACE}^L$.

By 12.3.19 there is a sentence φ_O in $\text{FO}(\text{DTC})[\text{rel-}\omega\text{-}\mathbf{Q}_L][\sigma]$ such that $O = \text{Mod}(\varphi_O)$ and hence, by 12.4.1, there is a log space-bounded oracle machine (M_O, L) that decides O .

We construct a log space-bounded oracle machine (M, L) that accepts K . Roughly, (M, L) runs as follows: Started with an ordered τ -structure \mathcal{A} , it imitates (M_K, O) . Suppose it reaches an oracle configuration of (M_K, O) . (M, L) does not have an oracle to decide whether the inscription J on the oracle tapes of (M_K, O) encodes a structure in O . However, using its oracle for L , it now can imitate a run of (M_O, L) to answer the question whether J encodes a structure in O . For this purpose, (M, L) must use work tapes that carry the inscription J to simulate the input tapes of (M_O, L) . But J , in general, is not log space-bounded in $\|\mathcal{A}\|$. To overcome this difficulty, one applies the same trick as for 7.5.23, going always back to obtain the information on J needed in a given situation.

To be more precise, consider a configuration c which is the starting configuration of (M_K, O) or the immediate successor of an oracle configuration of (M_K, O) . Assume that there is a next oracle configuration c' and that the transition from c to c' leads to the inscription J on the oracle tapes of (M_K, O) . Assume that c is stored on an additional tape. In order to decide whether J encodes a structure in O , the information about J that (M, L) really needs, is of the kind

“What letter stands in a certain cell of a certain oracle tape of (M_K, O) ?”

To provide such an information, (M, L) imitates (M_K, O) starting from c , neglecting the oracle tapes, but using counters to mark the positions of the heads on the oracle tapes. Finally, (M, L) replaces the stored c by c' . \square

We come back to the question (+) of the introduction. For FO(IFP) and PTIME and for FO(DTC) and LOGSPACE we give positive answers, thereby proceeding in two steps (Proposition 12.4.3 and Proposition 12.4.4).

Proposition 12.4.3 *Let $\langle \in \sigma$ and let O be a class of ordered σ -structures. Then we have:*

- (a) *Every class in PTIME^O is $\text{FO}(\text{IFP})(\text{rel-}\omega\text{-}\mathbf{Q}_O)$ -axiomatizable.*
- (b) *Every class in LOGSPACE^O is $\text{FO}(\text{DTC})(\text{rel-}\omega\text{-}\mathbf{Q}_O)$ -axiomatizable.*

In both (a) and (b) the axiomatizing sentence can be chosen without nested quantifiers from $\text{rel-}\omega\text{-}\mathbf{Q}_O$.

Proof. We first prove (b). Let $\langle \in \tau$ and $K \subseteq \mathcal{O}[\tau]$ be accepted by some log space-bounded oracle machine (M, O) . We have to show that K is axiomatizable in $\text{FO}(\text{DTC})(\text{rel-}\omega\text{-}\mathbf{Q}_O)$.

To give such an axiomatization, we proceed as in 7.3.7, using log space-bounded configurations of (M, O) that only refer to the input tapes and to the work tapes; but now we describe the transition from one oracle configuration to the next one. For this purpose we assume that the *oracle answer states* s', s'' in oracle instructions are distinct from each other and do not occur as successor states of other instructions. The main point here is to show that the predicates $\text{Oracle}^+(\bar{u}, \bar{v})$ and $\text{Oracle}^-(\bar{u}, \bar{v})$ as given below can be expressed in $\text{FO}(\text{DTC})(\text{rel-}\omega\text{-}\mathbf{Q}_O)$. $\text{Oracle}^{+/-}(\bar{u}, \bar{v})$ says:

- \bar{u} is the starting configuration of (M, O) or a configuration with an oracle answer state;
- \bar{v} is an oracle configuration of (M, O) ;
- \bar{v} is reached from \bar{u} without passing an oracle configuration in between;
- if (M, O) is started in \bar{u} with empty oracle tapes then, when reaching \bar{v} , the oracle tapes contain/do not contain the encoding of a structure in O .

The first three points do not cause difficulties; so we restrict ourselves to the last one. We consider the run of (M, O) that leads from \bar{u} to \bar{v} (\bar{u} and \bar{v} obeying the first three points above) as a run of a log space-bounded (τ, σ) -transducer. Then we can apply Theorem 7.5.23 (more precisely: the obvious generalization involving the parameters \bar{u}, \bar{v}) to get, with a fixed $e \geq 1$, formulas

$$(+) \quad \psi_{\text{uni}}(\bar{x}, \bar{u}, \bar{v}) \text{ and } \psi_R(\bar{x}_1, \dots, \bar{x}_r, \bar{u}, \bar{v}) \text{ for } R \in \sigma, R \text{ } r\text{-ary,}$$

in $\text{FO}(\text{DTC})[\tau]$ where \bar{x} and the \bar{x}_i have length e , and, with similar methods, a formula $\varphi_{\text{Str}}(\bar{u}, \bar{v})$ in $\text{FO}(\text{DTC})[\tau]$ such that $\varphi_{\text{Str}}(\bar{u}, \bar{v})$ says that

after the run from \bar{u} to \bar{v} the inscription on the oracle tapes encodes a σ -structure, namely $(\psi_{\text{uni}}(-, \bar{u}, \bar{v}), (\psi_R(-, \dots, -, \bar{u}, \bar{v}))_{R \in \sigma})$.

Then we see that

$$\psi(\bar{u}, \bar{v}) \quad := \quad Q_{\text{rel-}e\text{-}O}\bar{x}, \dots (\psi_{\text{uni}}(\bar{x}, \bar{u}, \bar{v}), (\psi_R(\bar{x}_R, \bar{u}, \bar{v}))_{R \in \sigma})$$

expresses that $(\psi_{\text{uni}}(-, \bar{u}, \bar{v}), (\psi_R(-, \dots, -, \bar{u}, \bar{v}))_{R \in \sigma}) \in O$. Therefore, we can describe the fourth point of $\text{Oracle}^+(\bar{u}, \bar{v})$ by

$$\varphi_{\text{Str}}(\bar{u}, \bar{v}) \wedge \psi(\bar{u}, \bar{v})$$

and the fourth point of $\text{Oracle}^-(\bar{u}, \bar{v})$ by $\neg\varphi_{\text{Str}}(\bar{u}, \bar{v}) \vee \neg\psi(\bar{u}, \bar{v})$.

The proof of (a) is similar, referring to a PTIME analogue of Theorem 7.5.23 (see Exercise 7.5.25).

The additional remark of this proposition is obvious. □

The preceding proof cannot be extended to PSPACE and FO(PFP). The reason is that the oracle tapes of a polynomially space-bounded oracle machine may get inscriptions of length exponential in the length of the input structure and hence, may not be definable by means of an interpretation. One way to preserve the validity consists in changing the oracle model, say, by suitably restricting the use of the oracle tapes.

We now turn to model checking:

Proposition 12.4.4 *Let $\langle \in \sigma$ and let O be a class of ordered σ -structures. Then we have:*

- (a) *Every FO(IFP)(rel- ω - \mathbf{Q}_O)-axiomatizable class of ordered structures is in PTIME^O (and hence, is decidable by a machine according to PTIME^O).*
- (b) *Every class of ordered structures axiomatizable by an FO(DTC)(rel- ω - \mathbf{Q}_O)-sentence without nested quantifiers of rel- ω - \mathbf{Q}_O is in LOGSPACE^O (and hence, is decidable by a machine according to LOGSPACE^O).*
- (c) *The assertion in (b) is true for all FO(DTC)(rel- ω - \mathbf{Q}_O)-axiomatizable classes of ordered structures if O is NPTIME-complete with respect to logspace reductions.*

Summing up:

- Corollary 12.4.5** (a) *For any class O of ordered σ -structures, FO(IFP)(rel- ω - \mathbf{Q}_O) captures PTIME^O .*
- (b) *If O is NPTIME-complete with respect to logspace reductions, then FO(DTC)(rel- ω - \mathbf{Q}_O) captures LOGSPACE^O . □*

Note that in (b), by 12.4.2, $\text{LOGSPACE}^O = \text{LOGSPACE}^{\text{NPTIME}}$, so FO(DTC)(rel- ω - \mathbf{Q}_O) captures $\text{LOGSPACE}^{\text{NPTIME}}$. For example, we have

$$\text{FO(DTC)(rel-}\omega\text{-}\mathbf{Q}_{\text{HAM}_{<}}) \text{ captures } \text{LOGSPACE}^{\text{NPTIME}}.$$

The preceding results yield normal form theorems for both the IFP and the DTC case: Over ordered structures, any FO(IFP)(rel- ω - \mathbf{Q}_O)-sentence is equivalent to a sentence without nested quantifiers from rel- ω - \mathbf{Q}_O . The same is true for FO(DTC)(rel- ω - \mathbf{Q}_O) provided O is NPTIME-complete with respect to logspace reductions.

Proof (of 12.4.4). (a) We proceed by induction on $\text{FO}(\text{IFP})(\text{rel-}\omega\text{-}\mathbf{Q}_O)$ -formulas and confine ourselves to the $\text{rel-}\omega\text{-}\mathbf{Q}_O$ -step. So let

$$\varphi = Q_{\text{rel-}k\text{-}O} \bar{x}, \bar{x}_1, \dots, \bar{x}_s [\varphi_{\text{uni}}(\bar{x}), \varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s)]$$

where $\sigma = \{R_1, \dots, R_s\}$ and $\varphi_{\text{uni}}, \varphi_1, \dots, \varphi_s$ are $\text{FO}(\text{IFP})(\text{rel-}\omega\text{-}\mathbf{Q}_0)[\tau]$ -formulas. By induction hypothesis there are polynomially time-bounded oracle machines $(M_{\text{uni}}, O), (M_1, O), \dots, (M_s, O)$ that decide

$$\{(\mathcal{A}, \bar{a}) \mid \mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \varphi_{\text{uni}}[\bar{a}]\}, \dots, \{(\mathcal{A}, \bar{a}) \mid (\mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \varphi_s[\bar{a}])\},$$

respectively. Using these machines, we can build a polynomially time-bounded oracle machine (M, O) deciding the class of ordered models of φ : Given an ordered τ -structure \mathcal{A} as input,

- (i) (M, O) checks whether $\varphi_{\text{uni}}^{\mathcal{A}}(-)$ is nonempty and whether (in case $R_1 = <$) $\varphi_1^{\mathcal{A}}(-)$ is an ordering on $\varphi_{\text{uni}}^{\mathcal{A}}(-)$; if not, it rejects, otherwise (viewing $(\varphi_{\text{uni}}, \varphi_1, \dots, \varphi_s)$ as an interpretation Π)
- (ii) using the ordering $\varphi_1^{\mathcal{A}}(-)$, it writes the (encoding of the) structure

$$\mathcal{A}^{\Pi} = (\varphi_{\text{uni}}^{\mathcal{A}}(-), \varphi_1^{\mathcal{A}}(-), \dots, \varphi_s^{\mathcal{A}}(-))$$

on some work tapes;

- (iii) it copies \mathcal{A}^{Π} on the oracle tapes;
- (iv) it invokes its oracle O to decide whether \mathcal{A}^{Π} belongs to O .

Note that in (ii) we cannot write \mathcal{A}^{Π} directly on the oracle tapes, because these tapes may be used by the machines $(M_{\text{uni}}, O), (M_1, O), \dots, (M_s, O)$ when constructing \mathcal{A}^{Π} .

(b) The proof follows the same pattern as that of (a). But now we write \mathcal{A}^{Π} directly on the oracle tapes; we can do so as the formulas $\varphi_{\text{uni}}, \varphi_1, \dots, \varphi_s$ do not contain quantifiers from $\text{rel-}\omega\text{-}\mathbf{Q}_O$. One only has to check that the arguments remain valid if the polynomial time bounds are replaced by logarithmic space bounds.

(c) The proof is more involved. The reason is that, when following the pattern of the proof of (a), we cannot realize (ii) by a log space-bounded oracle machine as the encoding of \mathcal{A}^{Π} may be of polynomial length, thus violating the logspace bound. We therefore proceed as follows: In the $\text{rel-}\omega\text{-}\mathbf{Q}_O$ -step of an inductive proof, let

$$\varphi = Q_{\text{rel-}k\text{-}O} \bar{x}, \bar{x}_1, \dots, \bar{x}_s [\varphi_{\text{uni}}(\bar{x}), \varphi_1(\bar{x}_1), \dots, \varphi_s(\bar{x}_s)]$$

where $\varphi_{\text{uni}}, \varphi_1, \dots, \varphi_s$ are $\text{FO}(\text{DTC})(\text{rel-}\omega\text{-}\mathbf{Q}_O)[\tau]$ -formulas. By hypothesis, there are log space-bounded oracle machines $(M_{\text{uni}}, O), \dots$ deciding $\{(\mathcal{A}, \bar{a}) \mid \mathcal{A} \in \mathcal{O}[\tau], \mathcal{A} \models \varphi_{\text{uni}}[\bar{a}]\}, \dots$, respectively. Using these machines, we now build a log space-bounded oracle transducer (M, O) that works as follows: Given an

ordered τ -structure \mathcal{A} as input, (M, O) checks whether $\varphi_{\text{uni}}^{\mathcal{A}}(-)$ is nonempty and whether $\varphi_1^{\mathcal{A}}(-)$ is an ordering on $\varphi_{\text{uni}}^{\mathcal{A}}(-)$; if not, it gives as output a fixed ordered σ -structure not in O (note that we may assume the existence of such a structure, the other case being trivial), otherwise it produces the output structure \mathcal{A}^{II} (cf. (ii) in the proof of (a)).

By construction of (M, O) we have:

$$\mathcal{A} \models \varphi \quad \text{iff} \quad (M, O), \text{ started with } \mathcal{A}, \text{ gives an output in } O,$$

i.e., the class of ordered models of φ is reducible to O via a log space-bounded oracle transducer. The next lemma tells us that we are done. \square

Lemma 12.4.6 *Let O be a class of ordered σ -structures that is NPTIME-complete with respect to logspace reductions and let K be a class of ordered τ -structures. Then we have:*

If K is reducible to O via a log space-bounded oracle transducer with oracle class \mathcal{O} , then K belongs to $\text{LOGSPACE}^{\mathcal{O}}$.

For the proof we need a kind of join of classes: Let σ and τ be disjoint and P a unary relation symbol that does not belong to $\sigma \cup \tau$. For $\mathcal{A} \in \text{Str}[\sigma]$ and $\mathcal{B} \in \text{Str}[\tau]$ let $\mathcal{A}\mathcal{B}$ be the $(\sigma \cup \tau \cup \{P\})$ -structure \mathcal{E} with (we assume that \mathcal{A} and \mathcal{B} are disjoint)

$$\begin{aligned} - E &:= A \dot{\cup} B \\ - P^{\mathcal{E}} &:= B \\ - R^{\mathcal{E}} &:= \begin{cases} R^{\mathcal{A}} & \text{if } R \in \sigma \\ R^{\mathcal{B}} & \text{if } R \in \tau. \end{cases} \end{aligned}$$

We extend the join to the case that one argument is the “empty structure” by setting

$$\begin{aligned} \mathcal{A}\emptyset &:= \text{the expansion of } \mathcal{A} \text{ to a } (\sigma \cup \tau \cup \{P\})\text{-structure } \mathcal{E} \\ &\quad \text{with } P^{\mathcal{E}} = \emptyset \text{ and } R^{\mathcal{E}} = \emptyset \text{ for all } R \in \tau; \\ \emptyset\mathcal{B} &:= \text{the expansion of } \mathcal{B} \text{ to a } (\sigma \cup \tau \cup \{P\})\text{-structure } \mathcal{E} \\ &\quad \text{with } P^{\mathcal{E}} = B \text{ and } R^{\mathcal{E}} = \emptyset \text{ for all } R \in \sigma. \end{aligned}$$

Note that $\mathcal{A}\mathcal{B} \cong \mathcal{A}'\mathcal{B}'$ implies $\mathcal{A} \cong \mathcal{A}'$ and $\mathcal{B} \cong \mathcal{B}'$. For $L \subseteq \text{Str}[\sigma]$ and $K \subseteq \text{Str}[\tau]$ let LK be the closure under isomorphisms of

$$\{\mathcal{A}\emptyset \mid \mathcal{A} \in L\} \cup \{\emptyset\mathcal{B} \mid \mathcal{B} \in K\}.$$

LK is a class of $(\sigma \cup \tau \cup \{P\})$ -structures. Up to trivial expansions it consists of the structures in L and K .

The proof of the following fact is routine.

Lemma 12.4.7 *If L and K are in NPTIME then $(LK)_{<}$ is in NPTIME.*

\square

Proof (of 12.4.6). Let O and K be given accordingly, and let K be reducible to O via the log space-bounded oracle transducer (M, O) . By 12.4.2 it suffices to show that $K \in \text{LOGSPACE}^{\text{NPTIME}}$, i.e., that there is a class O^* in NPTIME and a log space-bounded oracle machine (M^*, O^*) that accepts K .

Let (M, O) be started with an input structure of cardinality n . Thus, if n is sufficiently large, the part of the configurations of M comprising

- the current state
- the current inscriptions of the work tapes
- the current position of the heads on both the input and the work tapes

can be encoded by d numbers between 0 and $n - 1$ and hence, by d counters of length logarithmic in n . Here, d is a suitable positive number independent of n .

Let C_n be the set of these partial configurations. The cardinality of C_n is $\leq n^d$ and hence, polynomial in n . Given an input structure \mathcal{A} of cardinality n , set

$$\text{Or}(\mathcal{A}) := \{(c, \mathcal{B}) \mid c \in C_n \text{ and } (M, O), \text{ started with } \mathcal{A} \text{ and with empty oracle tapes in configuration } c, \text{ reaches a next oracle configuration, the oracle tapes then having an inscription that encodes } \mathcal{B}, \text{ and } \mathcal{B} \in O\}.$$

Now, considering sets $\{1, \dots, r\}$ as structures (of empty vocabulary), we let L consist of the structures $\mathcal{A}\{1, \dots, r\}$ with $r \in \mathbb{N}$ such that

- \mathcal{A} is an ordered τ -structure;
- there is a set $\{(c_1, \mathcal{B}_1), \dots, (c_r, \mathcal{B}_r)\} \subseteq \text{Or}(\mathcal{A})$ with distinct c_1, \dots, c_r such that $(M, \{\mathcal{B}_1, \dots, \mathcal{B}_r\})$, started with \mathcal{A} , gives an output in O .

Then we have:

- (1) For $\mathcal{A} \in \mathcal{O}[\tau]$ and $r = \|\text{Or}(\mathcal{A})\|$,

$$\mathcal{A}\{1, \dots, r\} \in L \quad \text{iff} \quad \mathcal{A} \in K.$$

- (2) $L_{<}$ is in NPTIME .

For (1) we note that $(M, \{\mathcal{B} \mid (c, \mathcal{B}) \in \text{Or}(\mathcal{A}) \text{ for some } c \in C_n\})$ started with \mathcal{A} behaves exactly like (M, O) started with \mathcal{A} .

To prove (2), one can define a (nondeterministic) machine that is polynomially time-bounded and, started with an input structure \mathcal{D} ,

- checks whether \mathcal{D} has the form $(\mathcal{A}\{1, \dots, r\}, <^{\mathcal{D}})$
- nondeterministically guesses distinct configurations $c_1, \dots, c_r \in C_{\|\mathcal{A}\|}$
- checks whether for all $i \in \{1, \dots, r\}$ there are \mathcal{B}_i such that $(c_i, \mathcal{B}_i) \in \text{Or}(\mathcal{A})$ (note that O is in NPTIME)
- runs $(M, \{\mathcal{B}_1, \dots, \mathcal{B}_r\})$ on \mathcal{A}

- accepts if the resulting output codes a structure in O (recall that O is in NPTIME).

Now we set

$$O^* := (OL)_<$$

which, by 12.4.7, belongs to NPTIME. Let σ^* be its vocabulary.

We are finished if we succeed in defining a log space-bounded (τ, σ^*) -oracle machine (M^*, O^*) that accepts K . We construct (M^*, O^*) in such a way that, started with an ordered τ -structure \mathcal{A} , it

- counts the number r_0 of configurations $c \in C_{\|\mathcal{A}\|}$ such that $(c, \mathcal{B}) \in \text{Or}(\mathcal{A})$ for some \mathcal{B} (to test this condition, (M^*, O^*) imitates (M, O) , starting with an ordered τ -structure \mathcal{A} in c ; if it reaches a next oracle configuration of (M, O) , the oracle tapes of (M^*, O^*) , that correspond to symbols of $\sigma^* \setminus \sigma$, are empty (and hence, code the empty relations); therefore, the inscription on the oracle tapes of (M, O) encodes a structure in O iff the inscription on the oracle tapes of (M^*, O^*) encodes a structure in O^* , and (M^*, O^*) can check this by use of its oracle);
- writes $\emptyset(\mathcal{A}\{1, \dots, r_0\})$ on its oracle tapes (more exactly, an ordered version, using the ordering of \mathcal{A});
- asks the oracle whether $\emptyset(\mathcal{A}\{1, \dots, r_0\}) \in O^*$ (or, equivalently, whether $(\mathcal{A}\{1, \dots, r_0\}) \in L$);
- accepts just in case the answer is positive.

By (1) we see that (M^*, O^*) accepts K . □

In 7.5.18 we have introduced the polynomial hierarchy PH. The following exercise contains an alternative definition. The complexity classes NPTIME^C appearing there are defined in a similar way as the classes LOGSPACE^C .

Exercise 12.4.8 Set $\text{PH}'_0 := \text{PTIME}$ and $\text{PH}'_k := \text{NPTIME}^{\text{PH}'_{k-1}}$ for $k \geq 1$. Show that $\text{PH}'_k = \text{PH}_k$ for all $k \geq 0$.

Notes 12.4.9 The notion of Lindström quantifier goes back to Lindström [118]; for further information see [31]. Various notions and results that we have treated in the book can be extended to logics with Lindström quantifiers; see, for example, [106, 27]. [89] discusses problems concerning the interplay between quantifiers and logical operators. The results of section 11.2 go back to [28] (Theorem 11.2.1), [80] (Theorem 11.2.3), and [16] (Theorem 11.2.5). For section 11.3, in particular Theorem 11.3.17, we refer to Dawar [24, 25]. Early references for reductions by interpretations are [120] and [22]. The results in section 11.4 are based on [51]; see also [53, 26]. Corollary 11.4.5(b) for $\text{HAM}_<$ goes back to Stewart [134].

References

1. S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Data Bases*. Addison-Wesley, 1995.
2. S. Abiteboul and V. Vianu. Fixpoint extensions of first-order logic and Datalog-like languages. In *Proc. 4th IEEE Symp. on Logic in Computer Science*, pages 71–79, 1989.
3. S. Abiteboul and V. Vianu. Datalog extensions for database queries and updates. *Journal of Computer and System Sciences* **43**, 62–124, 1991.
4. S. Abiteboul and V. Vianu. Generic computation and its complexity. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 209–219, 1991.
5. M. Ajtai and R. Fagin. Reachability is harder for directed than for undirected finite graphs. *Journal of Symbolic Logic* **55**, 113–150, 1990.
6. M. Ajtai and Y. Gurevich. Monotone versus positive. *Journal of the ACM* **34**, 1004–1015, 1987.
7. M. Ajtai and Y. Gurevich. DATALOG vs. first-order logic. In *Proc. 30th IEEE Symp. on Foundations of Computer Science*, pages 142–146, 1989.
8. J. Barwise. *Admissible Sets and Structures*. Springer, 1975.
9. J. Barwise. An introduction to first-order logic. In J. Barwise (ed.), *Handbook of Mathematical Logic*, pages 5–46. North Holland, 1977.
10. J. Barwise. On Moschovakis closure ordinals. *Journal of Symbolic Logic* **42**, 292–296, 1977.
11. Th. Behrendt, K. Compton, and E. Grädel. Optimization problems: Expressibility, approximation properties, and expected asymptotic growth of optimal solutions. In E. Börger et al. (eds.), *CSL '92, Lecture Notes in Computer Science* 702:43–60. Springer, 1993.
12. A. Blass, Y. Gurevich, and D. Kozen. A zero-one law for logic with a fixed-point operator. *Information and Control* **67**, 70–90, 1985.
13. E. Börger, E. Grädel, and Y. Gurevich. *The Classical Decision Problem*. Springer, 1997.
14. U. Bosse. An “Ehrenfeucht-Fraïssé Game” for fixpoint logic and stratified fixpoint logic. In E. Börger et al. (eds.), *CSL '92, Lecture Notes in Computer Science* 702:100–114. Springer, 1993.
15. J. R. Büchi. Weak second-order arithmetic and finite automata. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik* **6**, 66–92, 1960.

16. J. Cai, M. Fürer, and N. Immerman. An optimal lower bound on the number of variables for graph identification. *Combinatorica* **12**, 389–410, 1992. Extended abstract in *Proc. 30th Symp. on Foundations of Computer Science*, pages 612–617, 1989.
17. A. Calo and J.A. Makowsky. The Ehrenfeucht-Fraïssé game for transitive closure logic. In A. Nerode and M. Taitsin (eds.), *Logical Foundation of Computer Science - Tver '92, Lecture Notes in Computer Science* 620:57–68. Springer, 1992.
18. A. Chandra and D. Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences* **25**, 99–128, 1982. Extended abstract in *Proc. 21st Symp. on Foundations of Computer Science*, pages 333–347, 1980.
19. A. Chandra and D. Harel. Horn clause queries and generalizations. *Journal of Logic Programming* **1**, 1–15, 1985.
20. C. C. Chang and H. J. Keisler: *Model Theory*. North Holland, 3rd edition, 1990.
21. K. J. Compton. 0–1 laws in logic and combinatorics. In I. Rival (ed.), *NATO Advanced Study Institute on Algorithms and Order*, pages 353–383. D. Reidel, 1989.
22. E. Dahlhaus. Reductions to NP-complete problems by interpretations. In E. Börger et al. (eds.), *Logic and Machines: Decision Problems and Complexity, Lecture Notes in Computer Science* 171:357–365. Springer, 1983.
23. E. Dahlhaus. Skolem normal forms concerning the least fixpoint. In E. Börger (ed.), *Computation theory and Logic, Lecture Notes in Computer Science* 270:101–106. Springer, 1987.
24. A. Dawar. *Feasible Computation through Model Theory*. PhD thesis, University of Pennsylvania, 1993.
25. A. Dawar. Generalized quantifiers and logical reducibilities. *Journal of Logic and Computation* **5**, 213–226, 1995.
26. A. Dawar, G. Gottlob, and L. Hella. Capturing relativized complexity classes without order. *Mathematical Logic Quarterly* **44**, 109–122, 1998.
27. A. Dawar and E. Grädel. Generalized quantifiers and 0–1 laws. In *Proc. 10th IEEE Symp. on Logic in Computer Science*, pages 54–64, 1995.
28. A. Dawar and L. Hella. The expressive power of finitely many generalized quantifiers. In *Proc. 9th Symp. on Logic in Computer Science*, pages 20–29, 1994.
29. A. Dawar, L. Hella, and Ph. G. Kolaitis. Implicit definability and infinitary logic in finite model theory. In *Proceedings of the 22nd International Colloquium on Automata, Languages, and Programming (ICALP 95), Lecture Notes in Computer Science* 944:624–635. Springer, 1995.
30. A. Dawar, S. Lindell, and S. Weinstein. Infinitary logic and inductive definability over finite structures. *Information and Computation* **119**, 160–175, 1995.
31. H.-D. Ebbinghaus. Extended logics: The general framework. In J. Barwise and S. Feferman (eds.), *Model-Theoretic Logics*, pages 25–76. Springer, 1985.

32. H.-D. Ebbinghaus, J. Flum, and W. Thomas: *Mathematical Logic*. Springer, 2nd edition, 1994.
33. A. Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fundamenta Mathematicae* **49**, 129–141, 1961.
34. T. Eiter, G. Gottlob, and Y. Gurevich. Normal forms for second-order logic over finite structures, and classification of NP optimization problems. *Annals of Pure and Applied Logic* **78**, 111–125, 1996.
35. H. B. Enderton: *A Mathematical Introduction to Logic*. Academic Press, New York, 1972.
36. R. Fagin. Generalized first-order spectra and polynomial-time recognizable sets. In R. M. Karp (ed.), *Complexity of Computation, SIAM-AMS Proceedings* 7:43–73, 1974.
37. R. Fagin. Probabilities on finite models. *Journal of Symbolic Logic* **41**, 50–58, 1976.
38. R. Fagin. The number of finite relational structures. *Discrete Mathematics* **19**, 17–21, 1977.
39. R. Fagin, L. J. Stockmeyer, and M. Y. Vardi. A simple proof that connectivity separates existential and universal monadic second-order logics over finite structures. Research Report RJ 8647, IBM, 1992.
40. J. Flum. On bounded theories. In E. Börger et al. (eds.), *CSL '91, Lecture Notes in Computer Science* 626:111–118. Springer, 1992.
41. J. Flum. On the (infinite) model theory of fixed point logics. In X. Caicedo et al. (eds.), *Models, Algebras, and Proofs, Lecture Notes in Pure and Applied Mathematics* 203:67–75. Dekker, 1999.
42. J. Flum, M. Kubierschky, and B. Ludäscher. Total and partial well-founded datalog coincide. In E. Afrati and Ph. G. Kolaitis (eds), *Proc. 6th International Conference on Database Theory (ICDT), Lecture Notes in Computer Science* 1186:113–124. Springer, 1997.
43. R. Fraïssé. Sur quelques classifications des systèmes de relations (English summary). *Université d'Alger, Publications Scientifiques, Série A*, **1**, 35–182, 1954.
44. M. Fürer. The computational complexity of the unconstrained limited domino problem (with implications for logical decision problems). In *Logical Machines: Decision Problems and Complexity, Lecture Notes in Computer Science* 171:312–319. Springer, 1981.
45. H. Gaifman. On local and nonlocal properties. In J. Stern (ed.), *Logic Colloquium '81*, pages 105–135. North Holland, 1982.
46. M. R. Garey and D. S. Johnson: *Computers and Intractability - A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co., 1979.
47. Y. V. Glebskij, D. I. Kogan, M. I. Liogon'kij, and V. A. Talanov. Range and degree of realizability of formulas in the restricted predicate calculus. *Cybernetics* **5**, 142–154, 1969.
48. K. Gödel. Ein Spezialfall des Entscheidungsproblems der theoretischen Logik. *Ergebnisse eines mathematischen Kolloquiums* **2**, 27–28, 1929.

49. W. Goldfarb. The unsolvability of the Gödel class with identity. *Journal of Symbolic Logic* **49**, 1237–1252, 1984.
50. W. D. Goldfarb, Y. Gurevich, and S. Shelah. A decidable subclass of the minimal Gödel class with identity. *Journal of Symbolic Logic* **49**, 1253–1261, 1984.
51. G. Gottlob. Relativized logspace and generalized quantifiers over finite structures. *Journal of Symbolic Logic* **62**, 545–574, 1997.
52. G. Gottlob, N. Leone and H. Veith. Succinctness as a source of complexity in logical formalisms. *Annals of Pure and Applied Logic* **97**, 231–260, 1999.
53. E. Grädel. On logical descriptions of some concepts in structural complexity theory. In E. Börger et al. (eds.), *CSL '89, Lecture Notes in Computer Science* 440:163–175. Springer, 1990.
54. E. Grädel. On transitive closure logic. In E. Börger et al. (eds.), *CSL '91, Lecture Notes in Computer Science* 626:149–163. Springer, 1992.
55. E. Grädel, Ph. G. Kolaitis, and M. Y. Vardi. On the decision problem for two-variable first-order logic. *Bulletin of Symbolic Logic* **3**, 53–69, 1997.
56. E. Grädel and G. McCollm. Deterministic versus nondeterministic transitive closure logic. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, pages 58–63, 1992.
57. E. Grädel and M. Otto. Inductive definability with counting on finite structures. In E. Börger et al. (eds.), *CSL '92, Lecture Notes in Computer Science* 702:231–247. Springer, 1993.
58. E. Grädel and M. Otto. On logic with two variables. *Theoretical Computer Science*, to appear.
59. E. Grädel, M. Otto and E. Rosen. Two-variable logic with counting is decidable. In *Proc. 12th IEEE Symp. on Logic in Computer Science*, pages 306–317, 1997.
60. E. Grädel, M. Otto and E. Rosen. Undecidability results on two variable logics. *Archive for Mathematical Logic*, to appear. Extended abstract in *Proc. 14th Symp. on Theoretical Aspects of Computer Science STACS '97. Lecture Notes in Computer Science* 1200:249–260. Springer, 1997.
61. M. Grohe. *Fixpunktlogiken in der endlichen Modelltheorie*. Master's thesis, University of Freiburg i.Br., 1992.
62. M. Grohe. *The structure of fixed-point logics*. PhD thesis, University of Freiburg i.Br., 1994.
63. M. Grohe. Complete problems for fixed-point logics. *Journal of Symbolic Logic* **60**, 517–527, 1995.
64. M. Grohe. Arity hierarchies. *Annals of Pure and Applied Logic* **82**, 103–163, 1996.
65. M. Grohe. Fixed-point logics on planar graphs. In *Proceedings 13th IEEE Symp. on Logic in Computer Science*, pages 6–15, 1998.
66. M. Grohe. Finite variable logic in descriptive complexity. *Bulletin of Symbolic Logic* **4**, 345–398, 1998.

67. M. Grohe and L. Hella. A double arity hierarchy theorem for transitive closure logic. *Archive for Mathematical Logic* **35**, 157–171, 1996.
68. M. Grohe and J. Mariño. Definability and descriptive complexity on databases of bounded tree width. In *Proc. 7th Intern. Conference on Database Theory*, Jerusalem 1998.
69. M. Grohe and T. Schwentick. Locality of order-invariant first-order formulas. In *Proc. 23rd Intern. Symp. on Math. Foundations in Computer Science, Lecture Notes in Computer Science* 1450:75–95. Springer, 1998.
70. Y. Gurevich. Toward logic tailored for computational complexity. In M. M. Richter et al. (eds.), *Computation and Proof Theory, Lecture Notes in Mathematics* 1104:175–216. Springer, 1984.
71. Y. Gurevich. Logic and the challenge of computer science. In E. Börger (ed.), *Current Trends in Theoretical Computer Science*, pages 1–57. Computer Science Press, 1988.
72. Y. Gurevich. Zero-one laws. *Bulletin of the European Association for Theoretical Computer Science* **46**, 90–106, 1992.
73. Y. Gurevich. From invariants to canonization. *Bulletin of the European Association for Theoretical Computer Science* **63**, 115–119, 1997.
74. Y. Gurevich, N. Immerman, and S. Shelah. McColm’s conjecture. In *Proc. 9th IEEE Symp. on Logic in Computer Science*, pages 10–19, 1994.
75. Y. Gurevich and S. Shelah. Random models and the Gödel case of the decision problem. *Journal of Symbolic Logic* **48**, 1120–1124, 1983.
76. Y. Gurevich and S. Shelah. Fixed-point extensions of first-order logic. *Annals of Pure and Applied Logic* **32**, 265–280, 1986.
77. Y. Gurevich and S. Shelah. On finite rigid structures. *Journal of Symbolic Logic* **61**, 549–562, 1996.
78. P. Hájek. On logics of discovery. In G. Goos and J. Hartmanis (eds.), *Proc. 4th Conference on Mathematical Foundations of Computer Science, Lecture Notes in Computer Science* 32:100–114. Springer, 1977.
79. W. Hanf. Model-theoretic methods in the study of elementary logic. In J. Addison, L. Henkin, and A. Tarski (eds.), *The Theory of Models*, pages 132–145. North Holland, 1965.
80. L. Hella. Logical hierarchies in PTIME. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, pages 360–368, 1992.
81. L. Hella, Ph. G. Kolaitis, and K. Luosto. How to define a linear order on finite models. In *Proc. 9th IEEE Symp. on Logic in Computer Science*, pages 40–49, 1994.
82. L. Hella, Ph. G. Kolaitis, and K. Luosto. Almost everywhere equivalence of logics in finite model theory. *Bulletin of Symbolic Logic* **2**, 422–443, 1996.
83. L. Hella, L. Libkin, and Y. Nurmonen. Notions of locality and their logical characterizations over finite models. *Journal of Symbolic Logic*, to appear.
84. L. Henkin. Logical systems containing only a finite number of symbols. Les Presses de l’Université de Montreal. 48 pages, 1967.

85. J. Hintikka. Distributive normal forms in the calculus of predicates. *Acta Philosophica Fennica* **6**, 1–71, 1953.
86. I. Hodkinson. Finite variable logics. *Bulletin of the European Association for Theoretical Computer Science* **51**, 111–140, 1993.
87. J. Hopcroft and J. Ullman: *Introduction to Automata Theory, Languages and Computation*. Addison Wesley, 1979.
88. H. Imhof. Computational aspects of arity hierarchies. In: D. van Dalen and M. Bezem (eds.), *CSL '96, Lecture Notes in Computer Science* 1258:211–225. Springer, 1997.
89. H. Imhof. Fixed-point logic, generalized quantifiers, and oracles. *Journal of Logic and Computation* **7**, 405–425, 1997.
90. N. Immerman. Upper and lower bounds for first-order expressibility. *Journal of Computer and System Sciences* **25**, 76–98, 1982.
91. N. Immerman. Relational queries computable in polynomial time. *Information and Control* **68**, 86–104, 1986. Extended abstract in *Proc. 14th Annual ACM Symp. on Theory of Computing*, pages 147–152, 1982.
92. N. Immerman. Expressibility as a complexity measure: results and directions. In *Second Structure in Complexity Conference*, pages 194–202. IEEE Computer Society Press, 1987.
93. N. Immerman. Languages that capture complexity classes. *SIAM Journal on Computing* **16**, 760–778, 1987. Extended abstract in *Proc. 15th Annual ACM Symp. on Theory of Computing*, pages 347–354, 1983.
94. N. Immerman. Nondeterministic space is closed under complement. *SIAM Journal on Computing* **17**, 935–938, 1988.
95. N. Immerman. *Descriptive Complexity*. Springer, 1998.
96. N. Immerman and D. Kozen. Definability with a bounded number of bound variables. *Information and Computation* **83**, 121–139, 1989.
97. N. Immerman and E. S. Lander. Describing graphs: a first-order approach to graph canonization. In A. Selman (ed.), *Complexity Theory Retrospective*, pages 59–81. Springer, 1990.
98. M. Kaufmann. A counterexample to the 0–1 law for existential monadic second-order logic. Computational Logic Inc., Internal Note 32, 1987.
99. M. Kaufman and S. Shelah. On random models of finite power and monadic logic. *Discrete mathematics* **54**, 285–293, 1985.
100. H. J. Keisler: *Model Theory for Infinitary Logic*. North Holland, 1971.
101. Ph. G. Kolaitis. Implicit definability on finite structures and unambiguous computations. In *Proc. 5th IEEE Symp. on Logic in Computer Science*, pages 168–180, 1990.
102. Ph. G. Kolaitis. The expressive power of stratified logic programs. *Information and Computation* **90**, 50–66, 1991.
103. Ph. G. Kolaitis and M. N. Thakur. Polynomial-time optimization, parallel approximation, and fixedpoint logic. In *Proc. 8th Annual Structure in Complexity Conference*, pages 31–41, 1993.

104. Ph. G. Kolaitis and M. N. Thakur. Logical definability of NP-optimization problems. *Information and Computation* **115**, 321–353, 1994.
105. Ph. G. Kolaitis and M. N. Thakur. Approximation properties of NP minimization classes. *Journal of Computing and System Sciences* **50**, 391–411, 1995. Extended abstract in *Proc. 6th IEEE Symp. on Structure in Complexity Theory*, pages 353–366, 1991.
106. Ph. G. Kolaitis and J. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic* **74**, 23–75, 1995.
107. Ph. G. Kolaitis and M. Y. Vardi. The decision problem for the probabilities of higher order logics. In *Proc. 19th ACM Symp. on Theory of Computing*, pages 425–435, 1987.
108. Ph. G. Kolaitis and M. Y. Vardi. 0-1 laws and decision problems for fragments of second-order logic. *Information and Computation* **87**, 302–338, 1990. Extended abstract in *Proc. 3rd IEEE Symp. on Logic in Computer Science*, pages 2–11, 1988.
109. Ph. G. Kolaitis and M. Y. Vardi. Infinitary logic and 0–1 laws. *Information and Computation* **98**, 258–294, 1992. Extended abstract in *Proc. 5th IEEE Symp. on Logic in Computer Science*, pages 156–167, 1990.
110. Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite model theory. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, pages 46–57, 1992.
111. Ph. G. Kolaitis and M. Y. Vardi. 0–1 laws for fragments of second-order logic: an overview. In Y. N. Moschovakis (ed.), *Logic from Computer Science*, pages 265–286. Springer, 1992.
112. M. Kubierschky. *Remisfreie Spiele, Fixpunktlogiken und Normalformen*. Master's thesis, University of Freiburg i.Br., 1995.
113. M. Kubierschky. Yet another hierarchy theorem. *Journal of Symbolic Logic*, to appear.
114. J.-M. Le Bars. Fragments of existential second-order logic without 0–1 laws. In *Proc. 13th IEEE Symp. on Logic in Computer Science* 1998, to appear.
115. L. Libkin. On the forms of locality over finite models. In *Proc. 12th IEEE Symp. on Logic in Computer Science*, pages 204–215, 1997.
116. L. Libkin. On counting and local properties. In *Proc. 13th IEEE Symp. on Logic in Computer Science* 1998, to appear.
117. S. Lindell. An analysis of fixed point queries on binary trees. *Theoretical Computer Science* **85**, 75–95, 1991.
118. P. Lindström. First order predicate logic with generalized quantifiers. *Theoria* **32**, 186–195, 1966.
119. A. Livchak. The relational model for process controll. (Russian) *Automated Documentation and Mathematical Linguistics* **4**, 27–29, 1983.
120. L. Lovász and P. Gács. Some remarks on generalized spectra. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **23**, 547–554, 1977.

121. M. Mortimer. On languages with two variables. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **21**, 135–140, 1975.
122. Y. N. Moschovakis: *Elementary Induction on Abstract Structures*. North Holland, 1974.
123. W. Oberschelp. Asymptotic 0–1 laws in combinatorics. In D. Jungnickel (ed.), *Combinatorial theory, Lecture Notes in Mathematics* 969:276–292. Springer, 1982.
124. M. Otto. The expressive power of fixed-point logic with counting. *Journal of Symbolic Logic* **61**, 147–176, 1996.
125. M. Otto. *Bounded Variable Logics and Counting. A Study in Finite Models*. Lecture Notes in Logic 9. Springer, 1997.
126. L. Pacholski and W. Szwoast. The 0–1 law fails for the class of existential second-order Gödel sentences with equality. In *Proceedings 30th IEEE Symp. on Foundations of Computer Science*, pages 160–163, 1989.
127. L. Pacholski, W. Szwoast, and L. Tendera. Complexity of two-variable logic with counting. In *Proc. 12th IEEE Symp. on Logic in Computer Science*, pages 318–327, 1997.
128. C. H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
129. C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity. In *Proc. 20th ACM Symp. on Theory of Computing*, pages 229–234, 1988.
130. B. Poizat. Deux ou trois choses que je sais de L_n . *Journal of Symbolic Logic* **47**, 641–658, 1982.
131. D. Scott. A decision method for validity of sentences in two variables. *Journal of Symbolic Logic* **27**, 477, 1962.
132. S. Shelah and J. Spencer. Zero-one laws for sparse random graphs. *Journal of the AMS* **1**, 97–115, 1988.
133. J. R. Shoenfield: *Mathematical Logic*. Addison-Wesley, Reading, MA, 1967.
134. I. A. Stewart. Using the Hamiltonian path operator to capture NP. *Journal of Computer and System Sciences* **45**, 127–151, 1992.
135. A. Stolboushkin. Finitely monotone properties. In *Proc. 10th IEEE Symp. on Logic in Computer Science*, pages 324–330, 1995.
136. H. Straubing: *Finite Automata, Formal Logic, and Circuit Complexity*. Birkhäuser, 1994.
137. R. Szelepcsényi. The method of forced enumeration for nondeterministic automata. *Acta Informatica* **26**, 279–284, 1988.
138. W. Thomas. *Languages, Automata, and Logic*. Report 9607. Institute for Computer Science and Applied Mathematics. University of Kiel, 1996.
139. B. Trahtenbrot. The impossibility of an algorithm for the decision problem for finite domains. (Russian) *Doklady Akademii Nauk SSSR* **70**, 569–572, 1950. English translation: *American Mathematical Society, Translations, Series 2* **23**, 1–5, 1963.

140. B. Trahtenbrot. Finite automata and the logic of monadic predicates. (Russian). *Sibirskij Mat. Zhurnal* **3**, 103-131, 1962. Extended abstract in *Doklady Akademii Nauk SSSR* **140**, 326-329, 1961. English translation: *Soviet Physics. Doklady* **6**, 753-755, 1961.
141. J. Tyszkiewicz. On asymptotic probabilities of monadic second order properties. In E. Börger et al. (eds.), *CSL '91, Lecture Notes in Computer Science* 626:425-439. Springer, 1992.
142. M. Y. Vardi. The complexity of relational query languages. In *Proc. 14th ACM Symp. on Theory of Computing*, pages 137-146, 1982.
143. M. Y. Vardi. On the complexity of bounded-variable queries. In *Proc. 14th ACM Symp. on Principles of Data Base Systems*, pages 266-276, 1995.
144. M. Y. Vardi. What makes a modal logic so robustly decidable? In *Descriptive Complexity and Finite Models*. American Mathematical Society, 1997.
145. S. Vinner. A generalization of Ehrenfeucht's game and some applications. *Israel Journal of Mathematics* **12**, 279-298, 1972.
146. P. Winkler. Random structures and zero-one laws. In N. W. Sauer et al. (eds.), *Finite and infinite combinatorics in sets and logic, NATO Advanced Science Institutes Series*, pages 399-420. Kluwer, 1993.

Index

- τ, σ , 1
 R^A, c^A, R^A, c^A , 1
 $R^A a_1 \dots a_n$, 1
 $\{<, S, \min, \max\}$, 3
 τ_0 , 3
 $\mathcal{O}[\sigma]$, 3
 $\mathcal{A}|\tau$, 3
 $\mathcal{A} \cong \mathcal{B}$, 3
 $\mathcal{A} \times \mathcal{B}$, 3
 $\mathcal{A} \dot{\cup} \mathcal{B}, \mathcal{A} \triangleleft \mathcal{B}$, 4
 $\mathcal{A}^I, \dot{\cup}_I \mathcal{A}, \triangleleft_I \mathcal{A}$, 4
 $\|I\|$, 4
 \neg, \vee, \exists , 4
 $\wedge, \rightarrow, \leftrightarrow, \forall$, 5
 $\varphi(x_1, \dots, x_n)$, 6
 $\bar{x}, \varphi(\bar{x})$, 6
 $\mathcal{A} \models \varphi[a_1, \dots, a_n], \mathcal{A} \models \varphi$, 6
 $\Phi \models \psi, \models \psi$, 6
 $\Phi \models_{\text{fin}} \psi, \models_{\text{fin}} \psi$, 7
 $\bigwedge \Phi, \bigvee \Phi$, 7
 $\text{qr}(\varphi)$, 7
 $\Sigma_n, \Pi_n, \Delta_n$, 7
 $\varphi \frac{t_1 \dots t_n}{x_1 \dots x_n}, \varphi(t_1, \dots, t_n)$, 7
 $\exists \geq^n$, 7
 $\exists =^n, \exists \leq^n$, 8
 $\varphi \geq^n, \varphi =^n, \varphi \leq^n$, 8
 $\mathcal{A} \equiv \mathcal{B}$, 9
 $\varphi^A(-)$, 10
 $\mathcal{A} \equiv_m \mathcal{B}$, 14
 $p \subseteq q, \bar{a} \mapsto \bar{b}$, 15
 $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 16
 $G_m(\mathcal{A}, \mathcal{B})$, 17
 $\varphi_{\mathcal{A}, \bar{a}}^m, \varphi_{\mathcal{A}}^m, \varphi_{\bar{a}}^m(\bar{v})$, 18
 $W_m(\mathcal{A}, \mathcal{B})$, 20
 $\mathcal{A} \cong_m \mathcal{B}, (I_j)_{j \leq m} : \mathcal{A} \cong_m \mathcal{B}$, 20
 \mathcal{G}_I , 23
 $\mathcal{B}_I, \mathcal{D}_I$, 24
 $\text{IT}^m(\mathcal{A}, \bar{a})$, 26
 $S(r, a)$, 26
 $\mathcal{A} \equiv_m^{\text{MSO}} \mathcal{B}$, 38
 $\text{MSO-}G_m(\mathcal{A}, \mathcal{B})$, 38
 $\psi_{\bar{a}, \bar{p}}^j$, 38
 $(M)\Sigma_n^1, (M)\Pi_n^1$, 39
 Δ_n^1 , 40
 $L_{\infty\omega}, L_{\omega_1\omega}$, 40
 $\bigvee \Psi$, 40
 $\bigwedge \Psi$, 41
 $\mathcal{A} \equiv^{L_{\infty\omega}} \mathcal{B}$, 42
 $G_{\infty}(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 42
 $\mathcal{A} \cong_{\text{part}} \mathcal{B}, I : \mathcal{A} \cong_{\text{part}} \mathcal{B}$, 43
 $W_{\infty}(\mathcal{A}, \mathcal{B})$, 43
 T_{rand} , 45
 $L_{\infty\omega}^s, L_{\infty\omega}^{\omega}$, 47
 $\mathcal{A} \equiv^s \mathcal{B}, \mathcal{A} \equiv^{L_{\infty\omega}^s} \mathcal{B}, \mathcal{A} \equiv_m^s \mathcal{B}$, 48
 \bar{a}_i^s , 50
 $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}), G_{\infty}^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 50

- $G_m^s(\mathcal{A}, \mathcal{B}), G_\infty^s(\mathcal{A}, \mathcal{B}), 50$
 $\mathcal{A} \cong_m^s \mathcal{B}, (I_j)_{j \leq m} : \mathcal{A} \cong_m^s \mathcal{B}, 51$
 $\mathcal{A} \cong_{\text{part}}^s \mathcal{B}, I : \mathcal{A} \cong_{\text{part}}^s \mathcal{B}, 51$
 $\psi_{\bar{a}}^m, {}^s\psi_{\mathcal{A}, \bar{a}}, \psi_{\mathcal{A}}^m, 52$
 $W_m^s(\mathcal{A}, \mathcal{B}), W_\infty^s(\mathcal{A}, \mathcal{B}), 52$
 $\mathcal{A}/s, 54$
 $r(\mathcal{A}), 57$
 $\text{FO}(\mathcal{C}), L_{\infty\omega}(\mathcal{C}), C_{\infty\omega}, 59$
 $\exists \geq^l, 59$
 $\text{FO}(\mathcal{C})^s, C_{\infty\omega}^s, C_{\infty\omega}^\omega, 59$
 $\text{C-G}_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}), \text{C-G}_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b}), 60$
 $K <, 64$
 $L_n(K), L_n(\varphi), L_n(\tau), l_n(K), 71$
 $l(K), l_n(\varphi), l(\varphi), 72$
 $T_{\text{rand}}, 73$
 $l_n(K|H), l(K|H), 74$
 $T_{\text{rand}}(\varphi_0), 76$
 $U_n(K), u_n(K), u(K), 77$
 $u_n(K|H), 77$
 $u(K|H), 78$
 $\Sigma_1^1(\forall^* \exists^*), \Sigma_1^1(\exists^* \forall^*), 86$
 $\mathbb{A}, \mathbb{A}^*, \mathbb{A}^+, 105$
 $\lambda, 105$
 $\text{NDA}, \bar{\delta}, L(M), 106$
 $|u|, 107$
 $L^+, 108$
 $B_u, 109$
 $F_n, F_\infty, F^\varphi, 120$
 $\mathcal{L}_1 \leq \mathcal{L}_2, \mathcal{L}_1 \equiv \mathcal{L}_2, \mathcal{L}_1 < \mathcal{L}_2, 122$
 $\mathcal{L}[\tau], 122$
 $\text{TC}, \text{DTC}, 123$
 $s_0, s_+, s_-, 125, 133$
 $|j|_r, |j|_r^n, 131$
 $\text{co-}\mathcal{C}, 155$
 $\tau <, K <, 156$
 $\text{Str}[\tau], 158$
 $\text{PH}, 159$
 $[\text{IFP}_{\bar{x}, X} \varphi], 169$
 $[\text{GFP}_{\bar{x}, X} \psi(\bar{x}, X)] \bar{x}, 172$
 $F_{(\infty)}^i, 174$
 $\exists(\forall)x\psi(x), 179$
 $\delta_i^l(x_1, \dots, x_l, v, w), 179$
 $\varphi(Y \bar{y}), 179$
 $\#_x \varphi, 208$
 $\mathcal{A} \equiv_m^M \mathcal{B}, 213$
 $\text{MLFP-G}_m(\mathcal{A}, \mathcal{B}), 213$
 $\mathcal{A} \cong_m^M \mathcal{B}, 217$
 $\text{TC-G}_m^r(\mathcal{A}, \mathcal{B}), 229$
 $\gamma \leftarrow \gamma_1, \dots, \gamma_l, 239$
 $(\tau, \Pi)_{\text{int}}, (\tau, \Pi)_{\text{ext}}, 239$
 $\tau_{\text{int}}, \tau_{\text{ext}}, 239$
 $P_{(n)}^i, P_{(\infty)}^i, 240$
 $\mathcal{A}[\Pi], (\Pi, P)\bar{t}, 240$
 $\mathcal{A}[\Sigma], (\Sigma, P)\bar{t}, 241$
 $(\Pi, P)\bar{t}, 246$
 $(\Pi, Q)\bar{t}, 249$
 $\mathcal{L}[\tau]_0, \models^{\mathcal{L}, \tau}, 289$
 $\mathcal{L} \equiv_{\text{es}} \mathcal{C}, 290$
 $(M, f), 290$
 $\mathcal{L} \equiv_{\text{es}} \mathcal{C} \text{ on } S, 296$
 $\mathcal{A}^\Pi, \psi^{-\Pi}, 297$
 $Q_K \bar{x}_1, \dots, \bar{x}_s [\psi_1(\bar{x}_1), \dots, \psi_s(\bar{x}_s)], 308$
 $Q\bar{x}\psi(\bar{x}), 309$
 $\mathcal{L}(Q_K), \mathcal{L}(\mathbf{Q}), 309$
 $K^k, 312$
 $\mathbf{Q}_K^\omega, 312$
 $L_{\infty\omega}^\omega(\mathbf{Q}), 314$
 $K \leq_{\mathcal{L}} L, K \leq_{\mathcal{L}}^k L, 321$
 $\text{rel-}L, \text{rel-}k\text{-}L, 322$
 $Q_{\text{rel-}k\text{-}L}, \text{rel-}\omega\text{-}\mathbf{Q}_L, 322$
 $\mathcal{L}[\mathbf{Q}], 323$
 $(M, O), 331$
 $\mathcal{C}^K, \mathcal{C}^L, 331$
 $\mathcal{A}\mathcal{B}, 336$

- accept, 126, 132
- acceptable, 126
- acceptor, 162
- almost all, 72
- alphabet, 105
- antitone, 167, 172, 192
- approximable, 280
- APX, 281
- arity, 1
 - of a quantifier, 309
- arity hierarchy, 268, 272, 273
- assignment, 6
- ATC, 261
- automaton, 106
 - deterministic, 106
 - nondeterministic, 106
- axiomatizable
 - in a logic, 124
 - in FO, 14, 20, 150, 252
 - in FO(DTC)(rel- ω - \mathbf{Q}_O), 333, 334
 - in FO(IFP)(rel- ω - \mathbf{Q}_O), 333, 334
 - in FO(M-LFP), 217
 - in FO^s, 54
 - in $L_{\omega_1\omega}$, 41
 - in $L_{\infty\omega}^s$, 53
 - in $L_{\infty\omega}^w$, 53
 - in SO, 150
- back property, 20
 - *s*-, 51
- ball, 26
 - *r*-, 26
- base of an instruction, 132
- basic term, 248
- Beth property, 63
- Beth's Theorem, 63
- blank letter, 125, 130
- body of a clause, 239
- bounded, 57, 58
 - *s*-, 57
 - fixed-point, 204
- breadth
 - of a DATALOG program, 242
 - of an S-DATALOG program, 242
- canonization
 - \mathcal{C} -, 294
 - \mathcal{C} - on S , 301
 - LOGSPACE-, 294
 - PSPACE-, 294
 - PTIME-, 291–293
- capture, 157
 - a complexity class, 151, 288
 - effectively strongly, 290, 295, 325
 - effectively strongly on S , 296
 - PTIME effectively strongly, 290, 292, 294, 300
 - PTIME effectively strongly on GRAPH, 300
 - strongly, 157, 288
- cardinality, 4
- cell, 125
 - virtual, 125, 130
- circuit
 - Hamiltonian, 2, 113, 327
- class
 - bounded, 57
 - elementary, 14
 - elementary relative to, 14
 - finitely axiomatizable, 14
 - fixed-point bounded, 204
 - free parametric, 80
 - indiscernible, 222
 - nontrivial free parametric, 80
 - nontrivial parametric, 77
 - of finite structures, 14
 - of structures, 10
 - parametric, 75
 - *s*-bounded, 57
 - *s*-rigid, 204
- clause, 239
- clique, 2, 113
- closed subset, 222
- closed under isomorphisms, 10
- closure
 - alternating transitive, 261
 - deterministic transitive, 123, 220
 - plus, 108
 - positive, 108
 - transitive, 123, 220
- colour type, 61
- Compactness Theorem, 8
- compatible, 76
- complement of a complexity class, 155
- complete
 - NPTIME- with respect to FO(DTC) reductions, 328
 - NPTIME- with respect to logspace reductions, 327, 328
 - strongly PTIME- with respect to FO-reductions, 325
 - with respect to \mathcal{C} -reductions, 321
 - with respect to \mathcal{L} -reductions, 324
- Completeness Theorem, 8
- complexity class
 - as a class of ordered structures, 153

- in complexity theory, 153
- component
 - connected, 2, 30
 - of a graph, 2
- computable
 - PTIME-, 204, 291
- concatenation of languages, 107
- configuration, 128, 134
 - accepting, 134
 - oracle, 331
 - starting, 136
- CONN, 2, 23, 122
- connected component, 30
- connective, 4
- co-NPTIME, 157
- consequence, 6
- constant, 1
- constant symbol, 1
- cost, 276
- counting quantifier, 59, 309
- Craig property, 64
- Craig’s Theorem, 64
- cut, 275
- cycle, 2, 302

- database, 119
 - relational, 12
- DATALOG, 239
 - I-, 245
 - P-, 248
 - positive, 245, 250
 - S-, 241
 - stratified, 241
 - totally defined WF-, 265
 - well-founded, 264
 - WF-, 264
 - with counting, 248
- DATALOG \equiv E(LFP), 243, 245
- DATALOG program, 239
- DATALOG_r
 - S-, 242
- DATALOG(#), 248
- decidability
 - of validity for $\forall^2\exists^*$, 100
 - of validity for FO^2 , 98
- decidable, 126
- definable
 - explicitly, 63
 - implicitly, 63, 217
 - in first-order logic, 114
- degree of a point in a graph, 2
- depth, 170
- deterministic transitive closure, 123, 220
- diameter of a graph, 85
- digraph, 1
 - acyclic, 24, 30, 253
- distance function
 - in a graph, 2, 23
 - in an ordering, 22
- distinct, 75
- do(p), 15
- domain, 1
- DTC, 123, 220
- duplicator, 16

- E(LFP), 243
- edge, 1
- Ehrenfeucht’s Theorem, 18
- Ehrenfeucht-Fraïssé *see* game
- element
 - existential, 261
 - universal, 261
- enumerable, 126
- enumeration
 - effective of a complexity class, 295
- equality symbol, 4
- equivalence relation
 - invariant, 107, 111
- equivalent, 7
 - elementarily, 9
 - logically, 6
 - $L_{\infty\omega}$ -, 42
 - m -, 14
 - m - in MSO, 38, 39
 - on ordered structures, 152
- EVEN, 199, 209
- EVEN[τ], 21, 37, 41, 53, 59, 72, 156
- existence of a fixed-point, 166
- expression
 - plus free regular, 114
 - regular, 108
- expressive power of a logic, 122
- extension axiom, 45, 53, 72
- extensional, 239

- F (truth value), 7
- FALSE, 7, 10, 11, 241
- feasible solution, 276
- finite model property, 95
 - of $\forall^2\exists^*$, 100
 - of $\exists^*\forall^2\exists^*$, 103
 - of FO^2 , 96, 99
 - of modal logic, 99
- fixed-point, 120, 166
 - greatest, 167, 172
 - least, 167

- simultaneous, 174
- simultaneous least, 175
- fixed-point logic *see* logic, fixed-point
- FO, 4
- FO \equiv S-DATALOG₀, 243
- FO^s, 47
- FO(ATC), 261, 314
- FO(BFP), 235
- FO(BFP) < FO(LFP), 237
- FO(BFP) \equiv S-DATALOG, 242
- FO(BFP_r), 236
- FO(C), 59
- FO(C)^s, 59
- FO(DTC), 123, 327
- FO(DTC) and LOGSPACE, 142, 148, 151, 296, 304
- FO(DTC) \leq FO(IFP²), 273
- FO(DTC) \leq FO(LFP²), 271
- FO(DTC) \leq FO(LFP³), 270
- FO(DTC) < FO(PFP), 272
- FO(DTC) \leq FO(PFP¹), 271
- FO(DTC) \leq FO(PFP²), 269
- FO(DTC) $\not\equiv$ FO(TC), 223
- FO(DTC)(rel- ω - $\mathbf{Q}_{\text{HAM}<}$), 334
- FO(DTC)[rel- ω - $\mathbf{Q}_{\text{HAM}<}$], 328
- FO(DTC^r), 268
- FO(IFP), 121, 169
- FO(IFP) and FO(PFP), 154
- FO(IFP) and PTIME, 138, 149, 151, 297, 302
- FO(IFP) \equiv FO(S-IFP), 179
- FO(IFP) \equiv I-DATALOG, 246
- FO(IFP) \equiv Σ_1^1 , 155
- FO(IFP) \equiv SO, 155
- FO(IFP^r), 268
- FO(IFP, #), 208
- FO(IFP, #) and PTIME, 305, 306, 315
- FO(IFP, #) \equiv DATALOG(#), 248
- FO(IFP, #) \equiv FO(PFP_{PTIME}, #), 210
- FO(LFP), 170
- FO(LFP) \leq Δ_1^1 , 212
- FO(LFP) \equiv E(LFP), 244
- FO(LFP) \equiv FO(ATC), 261
- FO(LFP) \equiv FO(IFP), 173
- FO(LFP) \equiv FO(PFP), 203, 204
- FO(LFP) \equiv FO(PFP_{PTIME}), 205, 206
- FO(LFP) \equiv FO(S-LFP), 179
- FO(LFP) \leq IMP(FO), 219
- FO(LFP) \leq $L_{\infty\omega}^{\omega} \cap$ PTIME, 206
- FO(LFP) \equiv WF-DATALOG, 265
- FO(LFP^r), 268
- FO(LFP¹) < FO(IFP¹), 273
- FO(M-LFP), 212
- FO(PFP), 121, 191
- FO(PFP) and PSPACE, 137, 149, 151
- FO(PFP) \equiv FO(S-PFP), 193
- FO(PFP) \equiv IMP(FO(PFP)), 219
- FO(PFP) \leq $L_{\infty\omega}^{\omega}$, 199
- FO(PFP) \equiv P-DATALOG, 249
- FO(PFP^r), 268
- FO(PFP_{PTIME}), 205
- FO(PFP_{PTIME}, #), 210
- FO(PFP, #), 209
- FO(PFP, #) \leq $C_{\infty\omega}^{\omega}$, 209
- FO(3-PFP), 198
- FO(posDTC), 224
- FO(posDTC) \equiv FO(DTC), 224
- FO(posTC), 143, 159, 224
- FO(posTC) and NLOGSPACE, 143, 148
- FO(posTC) \equiv FO(TC), 226
- FO(S-IFP), 175
- FO(S-LFP), 176
- FO(S-PFP), 192
- FO(TC), 123, 159, 220
- FO(TC) and NLOGSPACE, 143, 151
- FO(TC) < FO(BFP), 237
- FO(TC) < FO(LFP), 231
- FO(TC) \equiv S-DATALOG₁, 243
- FO(TC^r), 221, 268
- FO(TC¹) \equiv MSO, 221
- forest, 3
 - <-, 59, 60
- formula
 - atomic, 5
 - bounded DATALOG, 250
 - bounded positive DATALOG, 250
 - DATALOG, 240
 - existential, 65
 - existential FO(posTC), 225
 - first-order, 4, 169
 - I-DATALOG, 246
 - $L_{\infty\omega}$, 40, 42
 - m -Hintikka, 18
 - monotone in a relation, 67
 - negative in X , 169
 - nontrivial, 254
 - normal, 95, 169
 - P-DATALOG, 249
 - positive DATALOG, 245
 - positive in a relation, 67
 - positive in X , 169
 - S-DATALOG, 241
 - s -Scott, 57
 - totally defined, 194
 - totally defined WF-DATALOG, 265

- universal, 65
- WF-DATALOG, 265
- forth property, 20
 - s -, 51
- FO[τ], 5
- frame, 99
- Fraïssé's Theorem, 21
- free parametric, 80
- free(φ), 5
- function
 - antitone, 167, 172, 192
 - inductive, 166, 175
 - inflationary, 120, 166, 175
 - monotone, 166, 175
 - polynomial time computable, 190
- function symbol, 11
- function system
 - inductive, 175
 - inflationary, 175
 - monotone, 175
- Gaifman graph, 26
- Gaifman's Theorem, 31
- game
 - associated with a digraph, 168, 172, 198
 - Ehrenfeucht-Fraïssé for counting quantifiers, 60
 - Ehrenfeucht-Fraïssé for FO, 16, 21
 - Ehrenfeucht-Fraïssé for FO(M-LFP), 213
 - Ehrenfeucht-Fraïssé for FO^s, 50
 - Ehrenfeucht-Fraïssé for FO(TC), 229
 - Ehrenfeucht-Fraïssé for L_{∞ω}, 42
 - Ehrenfeucht-Fraïssé for L_{∞ω}^s, 50
 - Ehrenfeucht-Fraïssé for MSO, 38
 - Ehrenfeucht-Fraïssé with infinitely many moves, 42
 - of life, 197
 - pebble, 50, 60
- global relation, 10, 12
 - L_{∞ω}^s-definable, 54
- graph, 1
- GRAPH, 1, 298, 322
- graph
 - acyclic, 2, 306
 - acyclic connected, 306
 - balanced, 112
 - bipartite, 112, 113, 221, 223
 - coloured, 61
 - connected, 2, 23, 29, 40, 41, 85, 86, 122, 123, 169, 221, 309
 - connected ordered, 23
 - directed, 1
 - Gaifman, 26
 - of a structure, 26
 - planar, 85
 - random, 85
 - rigid, 86
 - stable coloured, 61
 - undirected, 1
 - with a Hamiltonian circuit, 113
- grid, 302
- Halting Problem, 127
- HAM, 113, 327
- Hamiltonian circuit, 2, 113, 327
- Hamiltonian path, 2
- Hanf's Theorem, 27
- hard
 - with respect to \mathcal{C} -reductions, 321
 - with respect to \mathcal{L} -reductions, 324
- head
 - of a clause, 239
 - of a Turing machine, 125
 - read-and-write, 125, 130
 - read-only, 130
- HENK, 328
- hierarchy
 - arity, 268, 273
 - arity for FO(PFP), 272
 - polynomial, 159, 338
- Hintikka formula
 - m -, 18
- hold
 - almost surely, 72
 - in almost all finite structures, 72
- homomorphism, 34
 - strict, 34
- Horn sentence, 251
- IMP(FO) $\leq \Delta_1^1$, 219
- IMP(FO(PFP)) $\not\leq L_{\infty\omega}^\omega$, 219
- IMP(\mathcal{L}), 218
- in-degree, 2
- index of an equivalence relation, 107
- induction
 - on the length of formulas, 5
 - simultaneous, 178
 - simultaneous for LFP and IFP, 179
 - simultaneous for PFP, 193
- inductive, 166, 175
- inflationary, 120, 166, 175
- input inscription, 131
- input instance, 276
- input tape, 127, 130

- Instr(M), 125, 131
- instruction, 125, 131
 - oracle, 331
- intentional, 239
- interpolant, 64
- interpolation property, 64
- interpretation, 297, 321
 - of σ in τ , 297, 321
- invariant
 - s - of a structure, 55, 200
- invariantization
 - \mathcal{C} - on S , 301
 - PTIME-, 292–294, 301
 - PTIME- on GRAPH, 301
- isomorphic, 3
 - m -, 20
 - partially, 43
 - s - m -, 51
 - s -partially, 51
- isomorphism, 3
 - generalized partial, 213
 - partial, 15
 - s -partial, 50
- isomorphism type, 18, 77
 - m -, 18, 25, 26
 - s - m -, 52
- join
 - simultaneous, 178
- king, 97
- language, 106, 126
 - acceptable, 126
 - accepted by a Turing machine, 126
 - accepted by an NDA, 106
 - decidable, 126
 - decided by a Turing machine, 126
 - definable by a Σ_1^1 -sentence, 111
 - definable in FO, 114, 116
 - definable in monadic second-order logic, 109, 111
 - enumerable, 126
 - noncounting, 116
 - plus free regular, 114
 - recognized by an automaton, 111
 - recognized by an NDA, 106, 111
 - regular, 108, 110, 111
 - ultimately periodic, 112
- law
 - 0–1 for FO, 77
 - 0–1 for $L_{\infty\omega}^\omega$, 77
 - labeled 0–1, 72
 - labeled 0–1 for FO, 73
 - labeled 0–1 for FO(PFP), 200
 - labeled 0–1 for $L_{\infty\omega}^\omega$, 73
 - labeled 0–1 for parametric classes, 77
 - labeled 0–1 for Ψ , 77
 - labeled 0–1 for $\Sigma_1^1(\exists^*\forall^*)$, 88
 - unlabeled 0–1, 80
 - unlabeled 0–1, failure for Σ_1^1 -FO², 98
 - unlabeled 0–1 for FO, 80
 - unlabeled 0–1 for FO(PFP), 200
 - unlabeled 0–1 for $L_{\infty\omega}^\omega$, 80
 - unlabeled 0–1 for parametric classes, 80
 - unlabeled 0–1 for $\Sigma_1^1(\exists^*\forall^*)$, 88
- leaf, 3
- Lemma
 - Pumping, 107, 111
 - Simultaneous Induction, 178
 - Simultaneous Induction for LFP and IFP, 179
 - Simultaneous Induction for PFP, 193
 - Transitivity, 182
 - Transitivity for LFP, 182
 - Transitivity for PFP, 195
- length of a path, 2
- Lindström extension, 309
- Lindström quantifier, 309
 - vectorization of, 312
- log, 132
- log space-bounded, 160
- logarithmic space
 - deterministic, 132
 - nondeterministic, 132
- logic, 159, 289, 314
 - alternating transitive closure, 261, 314
 - bounded fixed-point, 235, 263
 - closed under order-invariant sentences, 64
 - deterministic transitive closure, 123
 - existential fixed-point, 243
 - first-order, 4, 13
 - first-order with counting quantifiers, 59
 - fixed-point with counting, 208, 248
 - inductive fixed-point, 169
 - infinitary, 40
 - inflationary fixed-point, 121, 169, 313
 - least fixed-point, 170, 263
 - monadic least fixed-point, 212
 - monadic second-order, 38, 110
 - partial fixed-point, 121, 191, 263, 313
 - partial fixed-point with counting, 209
 - positive existential fixed-point, 245

- second-order, 37, 210
- simultaneous inflationary fixed-point, 175
- simultaneous least fixed-point, 176
- simultaneous partial fixed-point, 192
- stratified fixed-point, 235
- transitive closure, 123, 220, 263
- logic program
 - general, 239
- logically equivalent, 6
- logically valid, 6
- LOGSPACE, 132
- LOGSPACE and FO(DTC), 142, 148, 151, 296, 304
- LOGSPACE^K, 331
- LOGSPACE^{NPTIME}, 332, 334
- LOGSPACE^{NPTIME} = LOGSPACE^{HAM<}, 332
- LOGSPACE^{NPTIME} = LOGSPACE^{HENK<}, 332
- LOGSPACE^O, 334
- Löwenheim-Skolem Theorem, 8

- MAXCLIQUE, 284
- MAXCUT, 275, 276, 284
- MAX FO, 278
- maximization, 276
- MAXIMUM CONNECTED COMPONENT, 285
- MAX PB, 278
- MAX Π_n , 278
- MAX Σ_n , 278
- MAX_e Σ_1 , 281
- MAX_e $\Sigma_1 \subseteq$ APX, 281
- method
 - Ehrenfeucht-Fraïssé, 21
- MIN FO, 278
- MIN PB, 278
- minimization, 276
- MINVERTEX, 276, 277
- Mod(Φ), 14
- Mod(φ), 10
- modal logic, 99
- model
 - minimal, 34, 35
 - of a formula, 6
 - random, 84
- model checker, 147
- model class, 10
- model theory, 26
- monotone, 166, 175
- move, 16, 168, 213, 229
 - \forall -, 25
 - \exists -, 25
 - in $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 50
 - LFP, 213
 - negative LFP, 213
 - point, 38, 213, 229
 - positive LFP, 213
 - set, 38
 - TC, 229
- MSO, 38, 212
- MSO- $G_m(\mathcal{A}, \mathcal{B})$, 38

- NDA, 106
- NLOGSPACE, 132
- NLOGSPACE and FO(posTC), 143, 148
- NLOGSPACE and FO(TC), 143, 151
- NLOGSPACE=co-NLOGSPACE?, 161
- noncounting, 116
- nontrivial parametric, 75
- normal form
 - for FO(IFP), 138, 152, 189, 253, 255
 - for FO(LFP), 174, 188, 254, 259
 - for FO(PFP), 138, 152, 192, 194, 196, 197, 255
 - for FO(posTC), 226
 - for FO(TC), 228
 - for FO²-sentences, 96
 - prenex, 7, 39
- not, 4
- NPSPACE, 126
- NPTIME, 126, 132, 328, 329
- NPTIME and Σ_1^1 , 139, 150, 151, 157
- NPTIME optimization problem, 276, 280
 - first-order definable, 277
 - polynomially bounded, 277
- NPTIME=co-NPTIME?, 158
- number part, 208
- number sort, 208
- number variable, 208
- numerical invariant, 204

- occurrence
 - bound, 5
 - free, 5, 6, 37, 41, 172
- operation
 - boolean, 111
- opt_Q, 276
- optimization problem
 - first-order definable, 277
 - NPTIME, 276, 280
 - polynomially bounded, 277
- or, 4

- oracle answer state, 333
- oracle configuration, 331
- oracle instruction, 331
- oracle machine
 - (τ, σ) , 330
- oracle tape, 330
- oracle transducer, 331
- ORD, 88
- order-invariant in the finite, 64, 156, 288
- ordered representation, 156
- ordering, 3, 88
 - as $\{<, S, \min, \max\}$ -structure, 3
 - lexicographic, 131
 - minimal, 219
 - of even cardinality, 22
- ordMod, 130
- out-degree, 2
- output tape, 162

- parametric, 75
- Part(A, B), 15
- partial isomorphism, 15
- partially isomorphic, 43
- path, 2
- pebble, 50, 60
- pebble game, 50, 60
- PFP(k, l), 263
- plus closure, 108
- plus free regular, 114
- point
 - drawn, 168, 172, 192, 198
 - of a graph, 1
 - won, 168, 172, 192, 198
- point sort, 208
- point variable, 208
- polynomial hierarchy, 159, 338
- polynomial space, 126, 132
 - nondeterministic, 126
- polynomial time, 126, 132
 - nondeterministic, 126, 132
- positive closure, 108
- Pow, 106
- predecessor of a vertex, 2
- prenex normal form, 7, 39
- preservation
 - under extensions, 66
 - under homomorphisms, 34
 - under strict homomorphisms, 34, 35
 - under substructures, 65
- preservation theorem, 65
- probability
 - labeled, 74
 - labeled asymptotic, 72, 74, 79, 80, 89
 - unlabeled, 77
 - unlabeled asymptotic, 77–80
- problem
 - maximization, 276
 - minimization, 276
- product of structures, 3, 24
- program
 - DATALOG, 239
 - general logic, 239
 - I-DATALOG, 245
 - P-DATALOG, 248
 - positive DATALOG, 250
 - S-DATALOG, 241
 - stratified DATALOG, 241
 - totally defined WF-DATALOG, 265
- programming language, 152
- projective Horn, 251–253
- proof
 - formal, 8
- PSPACE, 126, 132, 150
- PSPACE and FO(PFP), 137, 149, 151
- PTAS, 284
- PTIME, 126, 132, 261, 314, 315, 325
- PTIME and FO(IFP), 138, 149, 151, 297, 302
- PTIME and FO(IFP, #), 305, 306, 315
- PTIME^O, 334
- PTIME=NPTIME?, 155, 290
- PTIME=PSPACE?, 153, 154, 204, 206
- Pumping Lemma, 107, 111

- quantifier, 309
 - counting, 59, 309
 - existential, 4, 309
 - Lindström, 309
 - simple, 309
 - universal, 309
- quantifier rank, 7
 - for FO(M-LFP), 213
- query
 - PTIME-computable, 204
- query language, 12

- random structure
 - infinite, 45
- random structure theory, 45
- rank
 - s - of a structure, 57
- rank function, 185
- reducibility
 - many-one, 320
- reducible

- \mathcal{C} -, 320
- \mathcal{C} -many-one, 320
- \mathcal{L} -, 321
- logspace, 327
- reduct, 3
- refinement
 - stable coloured, 62
- regular expression, 108
- regular language, 108
- reject, 126
- relation, 1
 - defined in a structure, 10
 - global, 10, 12
 - global on a class, 10
 - nontrivial, 254
 - transitive closure, 11, 24
- relation symbol, 1
 - 0-ary, 241
 - extensional, 239
 - intentional, 239
- relativization
 - of a complexity class, 331
- represent effectively, 281
- $\text{rg}(p)$, 15
- RIG, 78, 83
- rigid, 78
- root, 3
- round, 168
- run, 125
 - accepting, 126
 - rejecting, 126

- satisfaction relation, 6
- satisfiable, 6
- satisfy in a structure, 6
- scattered
 - \mathcal{L} -, 30
- S-DATALOG \equiv S-DATALOG₂, 243
- second-order logic, 37
 - monadic, 38
- section, 177
- sentence, 5
 - $\forall^2\exists^*$ -, 100
 - basic local, 31
 - $\exists^*\forall^2\exists^*$ -, 103
 - existential positive, 34
 - first-order, 5
 - $L_{\infty\omega}$, 41
 - $L_{\omega_1\omega}$ -, positive in a relation, 67
 - local, 31
 - monadic Σ_1^1 -, 89
 - monotone in a relation, 67
 - nontrivial free parametric, 89
 - nontrivial parametric, 75, 86, 87
 - order-invariant in the finite, 64, 156, 288
 - parametric, 75
 - positive in a relation, 67
 - Σ_1^1 -, 110
 - $\Sigma_1^1(\forall^*\exists^*)$ -, 86
 - $\Sigma_1^1(\exists^*\forall^*)$ -, 86, 87
 - universal $L_{\omega_1\omega}$ -, 67
 - universal Horn, 251
- Σ_1^1 and NP-TIME, 139, 150, 151, 157
- simple, 309
- simultaneous join, 178
- SO, 37
- SO(PFP), 159
- SO \leq FO(PFP), 211
- space bound for oracle machines, 331
- space-bounded, 126
- SPACE(n^c), 272
- spectrum of a sentence, 46
- spoiler, 16
- square, 125
- stage, 166, 199
- stage comparison, 185
- start of a Turing machine, 125, 131
- starting configuration, 136
- state, 131
 - accepting, 106, 125, 131
 - final, 106
 - initial, 106, 125, 131
 - of an automaton, 106
 - of a Turing machine, 125
 - oracle answer, 333
 - rejecting, 125, 131
 - starting, 131
- State(M), 125, 131
- Str, 291
- Str $[\tau]$, 158
- string, 105
- strongly PTIME-complete, 325
- structure, 1
 - finite, 13
 - indiscernible, 222
 - infinite random, 45, 74
 - input, 276
 - labeled, 71
 - numerical, 161
 - of vocabulary τ , 1
 - ordered, 3, 130
 - random, 72
 - rigid, 78
 - s -rigid, 204
 - τ -, 1

- structures
 - $C_{\infty\omega}^s$ -equivalent, 61
 - elementarily equivalent, 9
 - equivalent in FO(M-LFP), 214
 - FO^s-equivalent, 48, 56
 - FO(TC)-equivalent, 230
 - $L_{\infty\omega}$ -equivalent, 42, 44
 - $L_{\infty\omega}^s$ -equivalent, 48, 52, 55, 57
 - m -equivalent, 14, 19, 21
 - m -equivalent in FO^s, 50, 52
 - m -equivalent in $L_{\infty\omega}^s$, 50
 - m -equivalent in MSO, 38
 - m -isomorphic, 21
 - s - m -isomorphic, 51, 52
 - s -partially isomorphic, 51, 52
 - partially isomorphic, 43, 44
- successor
 - of a configuration, 134
 - of a vertex, 2
- successor relation
 - k -adic, 172
- sum
 - ordered, of m -equivalent structures, 24, 25
 - ordered, of structures, 4, 24, 39
- supp(\bar{a}), 50
- support, 50, 83
- T (truth value), 7
- tape, 125
 - input, 127, 130
 - oracle, 330
 - output, 162
 - work, 127, 130
- τ -structure, 1
- TC, 11, 24, 123, 220
- term, 4
- Theorem
 - Beth’s, 63
 - compactness, 8
 - completeness, 8
 - Craig’s, 64
 - Ehrenfeucht’s, 18
 - Fraïssé’s, 21
 - Gaïfman’s, 31
 - Hanf’s, 27
 - Löwenheim-Skolem, 8
 - Trahtenbrot’s, 8, 127
- time-bounded, 126
- tournament, 75
- Trahtenbrot’s Theorem, 8, 127
- $T_{\text{rand}}(\varphi_0)$, 76, 84
- transducer
 - log space-bounded, 163
 - oracle, 331
 - polynomially time-bounded, 164
 - space-bounded, 162
 - (τ, σ) -, 162
- transitive closure logic *see* logic, transitive closure
- transition relation, 106
- transitive closure, 24, 123, 220
 - alternating, 261
- transitive closure relation, 11
- Transitivity Lemma, 182
 - for LFP, 182
 - for PFP, 195
- tree, 3, 170, 305
 - labeled binary, 206
- TREE, 3, 305
- tree-width, 171
- true, 6
- TRUE, 7, 10, 11, 241
- Turing machine, 125
 - deterministic, 125, 132
 - for τ -structures, 130
 - f space-bounded, 126, 132
 - f time-bounded, 126, 132
 - nondeterministic, 125, 132
 - (τ, σ) -, 162
 - with oracle tapes, 330
- type
 - m -isomorphism, 18, 25, 26
 - r -ball, 27
 - s - m -isomorphism, 52
- ultimately periodic, 112
- undecidability
 - of finite satisfiability for FO, 127
 - of order-invariance, 288
 - of satisfiability in finite graphs, 129
 - of the Halting Problem, 127
- union
 - disjoint, of structures, 4, 24, 39
 - of m -equivalent structures, 24
 - of structures, 4
- universal algebra, 26
- universe, 1
- valid
 - logically, 6
- variable, 4
 - free, 5, 6, 37, 41, 172
 - predicate, 37
 - relation, 37
- vectorization

- of a class, 312
- of a Lindström quantifier, 312
- relativized of a class, 322
- relativized of a Lindström quantifier, 322
- vertex, 1
- vertex cover, 275
- virtual letter, 125
- vocabulary, 1
 - containing function symbols, 11
 - function-free, 11
 - relational, 1
- width of an interpretation, 297, 321
- win
 - a pebble game, 50
 - $G_m(\mathcal{A}, \mathcal{B})$, 19, 21
 - $G_m(\mathcal{A}, \mathcal{B})$, 17
 - $G_m(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 18, 21
 - $G_\infty(\mathcal{A}, \mathcal{B})$, 44
 - $G_\infty(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 42, 43
 - $G_m^s(\mathcal{A}, \mathcal{B})$, 52
 - $G_m^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 52
 - $G_\infty^s(\mathcal{A}, \mathcal{B})$, 52
 - $G_\infty^s(\mathcal{A}, \bar{a}, \mathcal{B}, \bar{b})$, 52
 - $\text{MSO-}G_m((\mathcal{A}, \bar{P}, \bar{a}), (\mathcal{B}, \bar{Q}, \bar{b}))$, 38
- winning position, 20, 43, 52
- winning strategy, 17
- witness
 - negative, 263
 - positive, 263
 - strongly, 148, 153
- word, 105
 - empty, 105
- word model, 109, 221
- work tape, 127, 130