

Algorithmic Model Theory

SS 2016

Prof. Dr. Erich Grädel and Dr. Wied Pakusa

Mathematische Grundlagen der Informatik
RWTH Aachen



This work is licensed under:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

Dieses Werk ist lizenziert unter:

<http://creativecommons.org/licenses/by-nc-nd/3.0/de/>

© 2019 Mathematische Grundlagen der Informatik, RWTH Aachen.

<http://www.logic.rwth-aachen.de>

Contents

1	The classical decision problem	1
1.1	Basic notions on decidability	2
1.2	Trakhtenbrot's Theorem	7
1.3	Domino problems	13
1.4	Applications of the domino method	16
1.5	The finite model property	20
1.6	The two-variable fragment of FO	21
2	Descriptive Complexity	31
2.1	Logics Capturing Complexity Classes	31
2.2	Fagin's Theorem	33
2.3	Second Order Horn Logic on Ordered Structures	38
3	LFP and Infinitary Logics	43
3.1	Ordinals	43
3.2	Some Fixed-Point Theory	45
3.3	Least Fixed-Point Logic	48
3.4	Infinitary First-Order Logic	51
4	Expressive Power of First-Order Logic	55
4.1	Ehrenfeucht-Fraïssé Theorem	55
4.2	Hanf's technique	59
4.3	Gaifman's Theorem	61
4.4	Lower bound for the size of local sentences	66
5	Zero-one laws	73
5.1	Random graphs	73
5.2	Zero-one law for first-order logic	75
5.3	Generalised zero-one laws	80

6	Modal, Inflationary and Partial Fixed Points	85
6.1	The Modal μ -Calculus	85
6.2	Inflationary Fixed-Point Logic	87
6.3	Simultaneous Inductions	93
6.4	Partial Fixed-Point Logic	94
6.5	Capturing PTIME up to Bisimulation	97
7	Fixed-point logic with counting	105
7.1	Logics with Counting Terms	106
7.2	Fixed-Point Logic with Counting	107
7.3	The k -pebble bijection game	110
7.4	The construction of Cai, Fürer and Immerman	111

7 Fixed-point logic with counting

The (machine-independent) characterisation of complexity classes by logics (in the sense of Definition 2.4) yields deep insights into the structure of the classified problems. The theorem of Fagin (cf. Chapter 2) is a seminal result in the field of descriptive complexity theory, and gives such a correspondence between algorithmic and logical resources for the important class NP. If we restrict to ordered structures, we can also find such characterisation for PTIME as shown e.g. in the Immerman-Vardi theorem (cf. Chapter 3). However, it is still one of the major open questions in finite model theory whether there is a logic capturing PTIME on *all* finite structures. Note that if no such logic exists this would necessarily imply $\text{PTIME} \neq \exists\text{SO} = \text{NP}$.

As we will see, fixed-point logics, such as LFP or IFP, do not suffice to capture PTIME on arbitrary structures, and most of the naturally considered examples to separate them from PTIME involve some kind of counting. For instance, the simple class $\text{EVEN} = \{\mathfrak{A} : |A| \text{ is even}\}$ turns out to be not definable in LFP. Therefore Immerman proposed that counting quantifiers should be added to logics and asked whether a suitable variant of fixed-point logic with counting would suffice to capture PTIME.

Although Cai, Fürer and Immerman eventually answered this question negatively, the extension of fixed-point logic by counting terms (FPC) has turned out to be an important and robust logic, that defines a natural level of expressiveness. In this chapter we study the logic FPC and present the construction of Cai, Fürer and Immerman which yields the separation of FPC from PTIME. To be precise, we even present a slightly more general result which uses the concept of treewidth and which is due to Dawar and Richerby.

7.1 Logics with Counting Terms

There are different ways of adding counting mechanisms to a logic, which are not necessarily equivalent. The most straightforward possibility is the addition of quantifiers of the form $\exists^{\geq 2}$, $\exists^{\geq 3}$, etc., with the obvious meaning. While this is perfectly reasonable for bounded-variable fragments of first-order logic or infinitary logic it does not increase the expressiveness of logics such as FO or LFP, since they are closed under the replacement of $\exists^{\geq i}$ by i existential quantifiers. For fixed-point logic another severe restriction is that it does not allow for recursion over the counting parameters i in quantifiers $\exists^{\geq i}x$. These counting parameters should therefore be considered as variables that range over natural numbers. To define in a precise way a logic with counting and recursion, one extends the original objects of study, namely finite (one-sorted) structures \mathfrak{A} , to two-sorted auxiliary structures \mathfrak{A}^* with a second numerical (but also finite) sort.

Definition 7.1. With any one-sorted finite structure \mathfrak{A} with universe A , we associate the two-sorted structure $\mathfrak{A}^* := \mathfrak{A} \cup \langle \{0, \dots, |A|\}; \leq, 0, e \rangle$, where \leq is the canonical ordering on $\{0, \dots, |A|\}$, and 0 and e stand for the first and the last element. Thus, \mathfrak{A}^* is the disjoint union of \mathfrak{A} with a linear order of length $|A| + 1$.

For all logics we studied so far, we naturally obtain two-sorted variants defining properties of the extended structures \mathfrak{A}^* . For instance, formulas of two-sorted first-order logic over two-sorted vocabularies $\sigma \cup \{\leq, 0, e\}$ are evaluated in structures \mathfrak{A}^* where semantics are defined in the obvious way. From now on, we stick to the convention to use Latin letters x, y, z, \dots for the variables over the first sort, and Greek letters λ, μ, ν, \dots for variables over the second sort (the numerical sort). In counting logics, these two sorts are related by *counting terms*, defined by the following rule. Let $\varphi(x)$ be a formula with a variable x (over the first sort) among its free variables. Then $\#_x[\varphi]$ is a term in the second sort, with the set of free variables $\text{free}(\#_x[\varphi]) = \text{free}(\varphi) - \{x\}$. The value of $\#_x[\varphi]$ is the number of elements a that satisfy $\varphi(a)$.

We introduce counting logics starting with first-order logic with counting, denoted by FOC, which is the closure of two-sorted first-order

logic under counting terms. Here are two simple examples that illustrate the use of counting terms.

Example 7.2. On an undirected graph $G = (V, E)$, the formula $\forall x \forall y (\#_z [Exz] = \#_z [Eyz])$ expresses the assertion that every node has the same degree, i.e., that G is regular.

Example 7.3. We present below a formula $\psi(E_1, E_2) \in \text{FOC}$ which expresses the assertion that two equivalence relations E_1 and E_2 are isomorphic; of course a necessary and sufficient condition for this is that for every i , they have the same number of elements in equivalence classes of size i :

$$\psi(E_1, E_2) \equiv (\forall \mu) (\#_x [\#_y [E_1xy] = \mu] = \#_x [\#_y [E_2xy] = \mu]).$$

7.2 Fixed-Point Logic with Counting

We now define (*inflationary*) *fixed point logic with counting* (FPC) and *partial fixed point logic with counting* PFPC by adding to FOC the usual rules for building inflationary or partial fixed points, ranging over both sorts.

Definition 7.4. Inflationary fixed point logic with counting, FPC, is the closure of two-sorted first-order logic under the following rules:

- (1) The rule for building counting terms.
- (2) The usual rules of first-order logic for building terms and formulae.
- (3) The fixed-point formation rule. Suppose that $\psi(R, \bar{x}, \bar{\mu})$ is a formula of vocabulary $\tau \cup \{R\}$ where $\bar{x} = x_1, \dots, x_k$, $\bar{\mu} = \mu_1, \dots, \mu_\ell$, and R has mixed arity (k, ℓ) , and that (\bar{u}, \bar{v}) is a $k + \ell$ -tuple of first- and second-sort terms, respectively. Then

$$[\text{ifp } R\bar{x}\bar{\mu} . \psi](\bar{u}, \bar{v})$$

is a formula of vocabulary τ .

The semantics of $[\text{ifp } R\bar{x}\bar{\mu} . \psi]$ on \mathfrak{A}^* is defined in the same way as for the logic IFP, namely as the inflationary fixed point of the operator

$$F_\psi : R \longmapsto R \cup \{(\bar{a}, \bar{i}) \mid (\mathfrak{A}^*, R) \models \psi(\bar{a}, \bar{i})\}.$$

The definition of PFPC is analogous, where we replace inflationary fixed points by partial ones. In the literature, one also finds different variants of fixed-point logic with counting where the two sorts are related by *counting quantifiers* rather than counting terms. Counting quantifiers have the form $(\exists i x)$ for ‘there exist at least i different x ’, where i is a second-sort variable. It is obvious that the two definitions are equivalent. In fact, FPC is a very robust logic. For instance, its expressive power does not change if one permits counting over tuples, even of mixed type, i.e. terms of the form $\#_{\bar{x}, \bar{\mu}} \varphi$ (see exercise class). One can of course also define least fixed-point logic with counting, LFPC, but one has to be careful with the positivity requirement (which is more natural when one uses counting quantifiers rather than counting terms). The equivalence of LFP and IFP readily translates to LFPC \equiv IFPC.

Example 7.5. An interesting example of an FPC-definable query is the method of stable colourings for graph-canonization. Given a graph G with a colouring $f : V \rightarrow \{0, \dots, r\}$ of its vertices, we define a refinement f' of f , giving to a vertex x the new colour $f'x = (fx, n_1, \dots, n_r)$ where $n_i = \#y[Exy \wedge (fy = i)]$. The new colours can be sorted lexicographically so that they again form an initial subset of \mathbb{N} . Then the process can be iterated until a fixed point, the stable colouring of G is reached. It is easy to see that the stable colouring of a graph is polynomial-time computable and uniformly definable in FPC.

On many graphs, the stable colouring uniquely identifies each vertex, i.e. no two distinct vertices (i.e. vertices in different orbits of the automorphism group) get the same stable colour. In this way stable colourings provide a polynomial-time graph canonization algorithm for such classes of graphs. For instance, this is the case for the class of all trees or, more generally, any class of graphs with bounded treewidth.

We now discuss the expressive power and evaluation complexity of fixed-point logic with counting. We are mainly interested in FPC-formulae and PFPC-formulae without free variables over the second sort, so that we can compare them with the usual logics without counting.

Exercise 7.1. Even without making use of counting terms, IFP over two-sorted structures \mathfrak{A}^* is more expressive than IFP over \mathfrak{A} . To prove this, construct a two-sorted IFP-sentence ψ such that $\mathfrak{A}^* \models \psi$ if, and only if,

$|A|$ is even.

It is clear that counting terms can be computed in polynomial-time. Hence the data complexity remains in PTIME for FPC and in PSPACE for PFPC. We shall see below that these inclusions are strict.

Theorem 7.6. On finite structures,

(1) $\text{IFP} \subsetneq \text{FPC} \subsetneq \text{PTIME}$.

(2) $\text{PFP} \subsetneq \text{PFPC} \subsetneq \text{PSPACE}$.

7.2.1 Infinitary Logic with Counting

Let $C_{\infty\omega}^k$ be the infinitary logic with k variables $L_{\infty\omega}^k$, extended by the quantifiers $\exists^{\geq m}$ ('there exist at least m ') for all $m \in \mathbb{N}$. Further, let $C_{\infty\omega}^\omega := \bigcup_k C_{\infty\omega}^k$.

Proposition 7.7. $\text{PFPC} \subseteq C_{\infty\omega}^\omega$.

Due to the two-sorted framework, the proof of this result is a bit more involved than for the corresponding result without counting, but not really difficult (see exercise class).

The separation of FPC from PTIME has been established by Cai, Fürer, and Immerman. Their proof also provides an analysis of the method of stable colourings for graph canonization. We have described this method in its simplest form in Example 7.1. More sophisticated variants compute and refine colourings of k -tuples of vertices. This is called the *k-dimensional Weisfeiler–Lehman method* and, in logical terms, it amounts to labelling each k -tuple by its type in $k + 1$ -variable logic with counting quantifiers. It was conjectured that this method could provide a polynomial-time algorithm for graph isomorphism, at least for graphs of bounded degree. However, Cai, Fürer, and Immerman were able to construct two families $(G_n)_{n \in \mathbb{N}}$ and $(H_n)_{n \in \mathbb{N}}$ of graphs such that on one hand, G_n and H_n have $\mathcal{O}(n)$ nodes and degree three, and admit a linear-time canonization algorithm, but on the other hand, in first-order (or infinitary) logic with counting, $\Omega(n)$ variables are necessary to distinguish between G_n and H_n . In particular, this implies Theorem 7.6.

7.3 The k -pebble bijection game

In Chapter ?? we introduced Ehrenfeucht-Fraïssé games to characterize the equivalence of structures (or, to put it in another way, definability of classes) in first-order logic. More specifically, two relational structures \mathfrak{A} and \mathfrak{B} can be distinguished by an FO-sentence of quantifier-rank $\leq m$ if, and only if, Spoiler has a winning strategy in the m -move Ehrenfeucht-Fraïssé game played on \mathfrak{A} and \mathfrak{B} which was denoted by $EF_m(\mathfrak{A}, \mathfrak{B})$.

Our next aim is to introduce the k -pebble bijection game which is an extension of the standard Ehrenfeucht-Fraïssé game to capture definability in $C_{\infty\omega}^\omega$. We will use these games to show that a certain (polynomial-time decidable) class of graphs is not definable in $C_{\infty\omega}^\omega$. In particular, this yields the separation of FPC from PTIME by Proposition 7.7.

Definition 7.8. The k -pebble bijection game $k\text{-BG}(\mathfrak{A}, \mathfrak{B})$ is a two-player game played on relational structures \mathfrak{A} and \mathfrak{B} using k pairs of pebbles $(x_1, y_1), \dots, (x_n, y_n)$ that can be placed on pairs of elements $(a_1, b_1), \dots, (a_n, b_n) \in A \times B$ during a play. The goal of Player I, who is called *Spoiler*, is to show that $\mathfrak{A} \not\equiv_{C_{\infty\omega}^k} \mathfrak{B}$ while Player II, the *Duplicator*, claims that $\mathfrak{A} \equiv_{C_{\infty\omega}^k} \mathfrak{B}$.

A *position* in the game $k\text{-BG}(\mathfrak{A}, \mathfrak{B})$ is a (partial) assignment $(a_1, b_1), \dots, (a_n, b_n)$ of pebbles on $A \times B$, so formally, a position is a (partial) mapping $p : \{1, \dots, k\} \rightarrow A \times B$. The initial position is $p = \emptyset$.

At position p a play proceeds as follows: First, Spoiler selects a pair of pebbles $i \leq k$. Duplicator has to react with a bijection $h : A \rightarrow B$ which respects all remaining pairs of pebbled elements (except for i), i.e. for all $i \neq j \in \text{dom}(p)$ and $p(j) = (a_j, b_j)$ we have $h(a_j) = b_j$. Spoiler then chooses $a \in A$ and the position is updated to $(p|i \mapsto (a_i, b_i))$ where

$$(p|i \mapsto (a_i, b_i))(j) := \begin{cases} p(j) & j \neq i \\ (a, h(a)) & j = i. \end{cases}$$

Spoiler wins a play, if either $|A| \neq |B|$ (i.e. Duplicator cannot respond with a bijection), or the play eventually reaches a position p such that the induced mapping $p(\{1, \dots, k\})$ is not a partial isomorphism of \mathfrak{A} and \mathfrak{B} , i.e. if $p(\{1, \dots, k\}) \notin \text{Loc}(\mathfrak{A}, \mathfrak{B})$. Infinite plays are won by Duplicator.

Theorem 7.9. If Duplicator wins the game $k\text{-BG}(\mathfrak{A}, \mathfrak{B})$, then $\mathfrak{A} \equiv_{C_{\infty\omega}^k} \mathfrak{B}$.

Proof. We prove by induction that for all formulae $\varphi(x_1, \dots, x_k) \in C_{\infty\omega}^k$, structures \mathfrak{A} and \mathfrak{B} and all $a_1, \dots, a_k \in A$ and $b_1, \dots, b_k \in B$ we have that if $\mathfrak{A} \models \varphi(a_1, \dots, a_k)$ and $\mathfrak{B} \not\models \varphi(a_1, \dots, a_k)$ then Spoiler has a winning strategy for $k\text{-BG}(\mathfrak{A}, \mathfrak{B})$ starting from position $p(i) = (a_i, b_i)$.

The cases of quantifier-free formulae, Boolean connectivities and first-order quantifier follow as in the case of Ehrenfeucht-Fraïssé games (cf. lecture notes of mathematical logic). Hence, we only consider $\varphi = \exists^{\geq i} x_j \psi(x_1, \dots, x_k)$. For this case, a winning strategy for Spoiler can be defined in the following way:

- Spoiler selects the pair $j \leq k$.
- Duplicator reacts with a bijection $h : A \rightarrow B$ respecting the remaining pebbled pairs.

We set $X = \{a \in A : \mathfrak{A} \models \psi(a_1, \dots, a_n)\}$ and $Y = \{b \in B : \mathfrak{B} \models \psi(b_1, \dots, b_n)\}$. From the assumption we know that $|X| \geq i$ and $|Y| < i$, hence there is an $a \in X$ such that $h(a) \notin Y$. Spoiler selects the element a and the position is updated to $(p|j \mapsto (a_j, b_j))$. As we have $\mathfrak{A} \models \varphi(a_1, \dots, a_n)$ and $\mathfrak{B} \not\models \varphi(b_1, \dots, b_{j-1}, h(a), b_{j+1}, \dots, b_n)$ the claim follows by induction. Q.E.D.

We can use Theorem 7.9 to show that a class \mathcal{K} of finite structures is not definable in $C_{\infty\omega}^\omega$. In particular, note that $\mathcal{K} \notin C_{\infty\omega}^\omega$ also implies that $\mathcal{K} \notin \text{FPC}$ since we have $\text{FPC} \leq C_{\infty\omega}^\omega$.

Proposition 7.10. Let $(\mathfrak{A}_k)_{k \geq 1}$ and $(\mathfrak{B}_k)_{k \geq 1}$ be two sequences of structures such that for infinitely many k we have $\mathfrak{A}_k \in \mathcal{K}$, $\mathfrak{B}_k \notin \mathcal{K}$ and Duplicators wins $k\text{-BG}(\mathfrak{A}_k, \mathfrak{B}_k)$. Then \mathcal{K} cannot be defined in $C_{\infty\omega}^\omega$.

7.4 The construction of Cai, Fürer and Immerman

We now present the construction of Cai, Fürer and Immerman which yields the separation of FPC from PTIME. Throughout this section, let $G = (V, E)$ denote a connected graph with $\deg(v) \geq 2$ for all $v \in V$. Starting from G we define a family of graphs $(X_S(G))_{S \subseteq E}$ that result by replacing every vertex v in G by a gadget $Z(v)$ and interconnecting different gadgets according to the edge relation in G .

For every v we define the set of new vertices $Z(v)$ as

$$Z(v) := \{a_{vw}, b_{vw}, c_{vw}, d_{vw} : w \in vE\} \cup \{v^X : X \subseteq vE, |X| \text{ even}\}.$$

Vertices of the form a_{vw}, b_{vw} are called *outer vertices* and they are intended to connect the two gadgets $Z(v)$ and $Z(w)$. The vertices c_{vw}, d_{vw} are *colour vertices* which are used only to make the set of outer nodes first-order definable. The remaining vertices v^S are called the *inner vertices*.

Let $X_{\emptyset}(G)$ denote the graph over the vertex set $\bigcup_{v \in V} Z(v)$ with the following edges:

- $(a_{vw}, c_{vw}), (b_{vw}, c_{vw}), (d_{vw}, c_{vw})$ for $(v, w) \in E$,
- (a_{vw}, v^X) for $w \in X$,
- (b_{vw}, v^X) for $w \notin X$, and
- (a_{vw}, a_{wv}) and (b_{vw}, b_{wv}) for all $(v, w) \in E$.



Figure 7.1. Example: gadget for a vertex v of degree three

In Figure 7.1 the construction of a gadget $Z(v)$ is illustrated for the case of a vertex v with degree three. The pairs of outer nodes a_{vx}, b_{vx} , a_{vy}, b_{vy} and a_{vz}, b_{vz} are connected to the corresponding outer nodes of the gadgets $Z(x)$, $Z(y)$ and $Z(z)$, respectively (this is indicated by the dashed lines in the figure).

We now extend the construction: for any (symmetric) set $S \subseteq E$ we define $X_S(G)$ to be the graph $X_{\emptyset}(G)$ in which for all $(v, w) \in S$ the edges (a_{vw}, a_{wv}) and (b_{vw}, b_{wv}) are replaced by (a_{vw}, b_{wv}) and (a_{wv}, b_{vw}) . We say that the edges in S have been *twisted*. In this way we obtain for every subset $S \subseteq E$ of edges a *CFI-graph* $X_S(G)$. Interestingly, we are going to show that these CFI-graphs $X_S(G)$ are completely determined by the parity of the set S :

Lemma 7.11. For all $S, T \subseteq E$ we have:

$$X_S(G) \cong X_T(G) \Leftrightarrow |S| \equiv |T| \pmod{2}.$$

Before we prove this claim in general, we consider some special cases. First of all, let all twisted edges be incident with a single vertex v .

Lemma 7.12. Let $S, T \subseteq vE$ be sets of neighbours of some vertex $v \in V$. If $S\Delta T = (S \setminus T) \cup (T \setminus S)$ is even, then

$$X_{v \times S}(G) \cong X_{v \times T}(G).$$

Proof. The mapping $\pi_{v;S;T} : X_{v \times S}(G) \rightarrow X_{v \times T}(G)$ defined by

$$\pi_{v;S;T}(z) := \begin{cases} z, & z \notin Z(v) \text{ or } z \text{ colour vertex,} \\ z, & z \in \{a_{vw}, b_{vw}\}, (v, w) \in S \cap T, \\ b_{vw}, & z = a_{vw}, (v, w) \in S\Delta T, \\ a_{vw}, & z = b_{vw}, (v, w) \in S\Delta T, \\ v^{X\Delta(S\Delta T)}, & z = v^X, \end{cases}$$

is an isomorphism (use that since X and $S\Delta T$ are even, the same holds for the symmetric difference $X\Delta(S\Delta T)$). Q.E.D.

We proceed to explain how one obtains an isomorphism between $X_{\{e\}}(G)$ and $X_{\{f\}}(G)$ for two distinct edges e and f of G .

Lemma 7.13. $X_{\{e\}}(G) \cong X_{\{f\}}(G)$.

Proof. If e and f are incident with the same vertex v , then the claim follows by Lemma 7.12. Hence, let $e = (u, v)$ and $f = (x, y)$ be such that $\{u, v\} \cap \{x, y\} = \emptyset$. Choose a path $v = v_1, v_2, \dots, v_\ell = x$ connecting v and x with $v_i \notin \{u, y\}$ for all $i \geq 1$. Then

$$\pi_{e \rightarrow f} := \pi_{v_1;u;v_2} \circ \pi_{v_2;v_1;v_3} \circ \dots \circ \pi_{v_{\ell-1};v_{\ell-2};x} \circ \pi_{v_\ell;v_{\ell-1};y},$$

is an isomorphism of $X_{\{e\}}(G) \cong X_{\{f\}}(G)$: the twist at edge (u, v) is moved along the path to the twist at edge (x, y) where both twists cancel out each other. Note that along the path, at every inner node v_i we

have precisely two twists of edges for the gadget $Z(v_i)$ which, again by Lemma 7.12, preserves the structure of the inner nodes. Q.E.D.

We are now ready to prove Lemma 7.11.

Proof (of Lemma 7.11). First of all, let $|S| \equiv |T| \pmod{2}$. If $|S| = |T| = 1$, then the claim follows by Lemma 7.13, so assume that $|S| \geq 2$ (or analogously, $|T| \geq 2$). Choose $e, f \in S$ with $e \neq f$. If e and f are incident with the same vertex $v \in V$ we know that $X_{S \setminus \{e, f\}}(G) \cong X_S(G)$ by Lemma 7.12. In the other case, we use the isomorphism $\pi_{e \rightarrow f}$ and see that $X_{S \setminus \{e, f\}}(G) \cong X_S(G)$. The claim follows by induction on $|\Delta T|$.

For the other direction assume that $\pi : X_{\{f=(x,y)\}}(G) \rightarrow X_{\emptyset}(G)$ is an isomorphism. Clearly, π maps outer (inner, colour) nodes to outer (inner, colour) nodes, and since π also induces an isomorphism of G , we can assume that for all $v \in V$ we have $\pi(Z(v)) = Z(v)$ and $\pi(\{a_{vw}, b_{vw}\}) = \{a_{vw}, b_{vw}\}$ for all $(v, w) \in E$. At this point we observe that if π interchanges a_{vw} and b_{vw} it necessarily interchanges a_{vw} and b_{wv} for all edges $(v, w) \in E$ except for (x, y) . Hence, the total number of interchanges of a 's and b 's in π is odd. This contradicts Lemma 7.12, however, as the number of interchanges of a 's and b 's in π for each gadget has to be even. Q.E.D.

We conclude that, up to isomorphism, there are precisely two CFI-graphs for G and we fix two canonical representatives from the isomorphism classes:

- $X(G) := X_{\emptyset}(G)$ (the *even CFI-graph* for G)
- $\tilde{X}(G) := X_{\{e\}}(G)$ for some edge $e \in E$ (the *odd CFI-graph* for G)

The *CFI-query* is to decide, given a CFI-graph $X_S(G)$, whether $X_S(G)$ is even or odd, i.e. whether $X_S(G) \cong X(G)$ or $X_S(G) \cong \tilde{X}(G)$.

Theorem 7.14. The CFI-query can be decided in polynomial time.

Proof. In order to count the number of twists, we need to identify the a and b -vertices. To this end it suffices to fix in every gadget $Z(v)$ an arbitrary inner node and to associate the intended labeling to the gadget $Z(v)$ (e.g. declare this node to be v^{\emptyset} and assign to all connected vertices b -labels and to the remaining outer ones a -labels). Then it is straightforward to count the number of twists modulo two. Lemma 7.11

guarantees that the isomorphism class of the resulting $\{a, b\}$ -labeled graph is independent of the initial choice of inner vertices. Q.E.D.

We conclude that the even and odd CFI-graphs can be distinguished in polynomial time. However, we are going to show that they cannot be separated by sentences in $C_{\infty\omega}^\omega$ if we start from a class of graphs G with sufficient complexity. In order to measure the complexity of graphs we introduce the important and well-studied concept of *treewidth*. Intuitively the treewidth of a graph formalises to what extent an (undirected) graph resembles a tree, and one of the reasons for its importance is that many NP-hard problems (and even some PSPACE-hard ones) become tractable on classes of graphs with bounded treewidth. There are various equivalent ways to characterize the treewidth of a graph, of which we sketch two: an algebraic and a game theoretic approach.

Definition 7.15. Let $G = (V, E)$ be an undirected graph. A *tree decomposition* of G is an undirected tree $\mathcal{T} = (T, E_T)$ where T is a family of subsets of V , i.e. $T \subseteq \mathcal{P}(V)$ and

- (a) $\bigcup T = V$, and
- (b) for all $(u, v) \in E$ there is some $X \in T$ so that $\{u, v\} \subseteq X$, and
- (c) for every vertex $v \in V$ the set $\{X \in T : v \in X\}$ is connected in \mathcal{T} .

Nodes in the tree \mathcal{T} are called *bags*. The *width* of the tree decomposition $\mathcal{T} = (T, E_T)$ is $(\max\{|X| : X \in T\} - 1)$, and the *treewidth* of G , denoted by $\text{tw}(G)$, is defined to be the minimal width for which a tree decomposition of G exists.

Next, we describe a game which characterises the notion of treewidth. The *k-cops and robber game* on G is played by two players, Player I (the cops) and Player II (the robber). The rules are as follows: the cops possess k pebbles (cops) which they can place on vertices of the graph. The robber has one pebble which is moved along paths. In each move the cops first choose some pebble which is either currently not placed on a vertex of the graph or which is removed from its current position w . Secondly, the cops determine a vertex v to be the new position for this pebble. After that, the robber reacts by moving his pebble along a path to a new vertex (which may be the old one). The chosen path has to be cop-free where the vertices v and w count as cop-free for

the current move. The cops win a play if, and only if, they can reach a position such that the robber cannot move. All other plays, i.e. all infinite ones, are won by the robber.

Seymour proved that a graph G has treewidth k if, and only if, the cops have a winning strategy in the game with $k + 1$ pebbles, but the robber wins the game if the cops are restricted to k pebbles. We use this game-theoretic characterisation to show:

Theorem 7.16. Let $G = (V, E)$ be a graph with $\delta(G) \geq 2$ and $\text{tw}(G) \geq k$. Then

$$X(G) \equiv_{\text{C}_{\infty}^k} \tilde{X}(G).$$

Proof. For two vertices u, v let $\sigma[u, v]$ be the permutation which exchanges u and v and fixes all other points. We say that a bijection $h : X(G) \rightarrow \tilde{X}(G)$ is *good except at node* $u \in V$ if

- $h(Z(v)) = Z(v)$ for all $v \in V$,
- h maps inner vertices to inner vertices and outer vertices to outer vertices,
- h is an isomorphism between the subgraphs $X(G) \setminus \{v^X : X \subseteq vE\}$ and $\tilde{X}(G) \setminus \{v^{\tilde{X}} : X \subseteq vE\}$, and
- for every pair $(a_{uv}, b_{uv}) \in Z(u)$, the mapping $h \circ \sigma[a_{uv}, b_{uv}]$ is an isomorphism from $X(G)[Z(u)]$ to $\tilde{X}(G)[Z(u)]$.

Let $\tilde{X}(G) = X_{(u,v)}(G)$. Then for instance $\sigma[a_{uv}, b_{uv}]$ is good except at u and $\sigma[a_{vu}, b_{vu}]$ is good except at v . Note that if $\eta \in \text{Aut}(\tilde{X}(G))$ with $\eta(Z(v)) = Z(v)$ for all $v \in V$ and h is good except at vertex u , then $h \circ \eta$ is good except at u as well.

The property of being good at some vertex can be propagated along a path in G : let P be a simple path in G from u to v , $P : u = v_1, v_2, \dots, v_{l-1}, v_l = v$, and let h be a bijection which is good except at vertex u . Then the bijection $h' := h \circ \eta_P$ where

$$\eta_P := \sigma[a_{uv_2}, b_{uv_2}] \circ \pi_{v_2v_1;v_3} \circ \dots \circ \pi_{v_{l-1}v_{l-2};v_l} \circ \sigma[a_{vv_{l-1}}, b_{vv_{l-1}}],$$

is good except at v and for $w \notin P, x \in Z(w)$ we have $h'(x) = h(x)$.

Finally, we describe a winning strategy for Duplicator in the k -

pebble bijection game played on $X(G)$ and $\tilde{X}(G)$. The strategy satisfies that pairs of pebbles (a_i, b_i) are always placed on vertices in a common gadget $Z(v)$. First of all, we initialize an instance of the k -cops and robber game played on G where we identify each of the k pairs of pebbles with one of the cops, and we assume that the robber makes his moves according to a fixed winning strategy (recall that $\text{tw}(G) \geq k$). The positions in the two games are related as follows: the vertex in G occupied by the i -th cop is precisely the vertex $v \in V$ for which the corresponding gadget $Z(v)$ in $X(G)$ and $\tilde{X}(G)$ is pebbled with the i -th pair (a_i, b_i) of pebbles in the k -pebble bijection game. We update the positions in the cops and robber game after each round of the k -pebble bijection game accordingly. Furthermore, whenever the robber is at some vertex $v \in V$, then Duplicator chooses in her current move some bijection which is good except at vertex v . For convenience, we assume that the robber starts at node u , and that in the first round Duplicator answers with the bijection $\sigma[a_{uv}, b_{uv}]$. Recall that this bijection is good except at vertex u .

We proceed to show that Duplicator can maintain the following invariant during each play: let $((a_1, \dots, a_k), (b_1, \dots, b_k))$ be the current position in the k -pebble bijection game, then

there is a bijection $g : X(G) \rightarrow \tilde{X}(G)$ with $g(a_i) = b_i$ for $i \leq k$ such that g is good except at a vertex $u \in V$ and for $i \leq k$ we have $a_i, b_i \notin Z(u)$ (u is the robber's position in the cops and robber game).

This can be seen as follows: assume Spoiler chooses the i -th pair of pebbles. Duplicator answers with the bijection g and Spoiler puts the i -th pair of pebbles onto some tuple $(a, g(a))$. By the condition on g of being good except at u , the new position in the k -pebble bijection game is indeed a partial isomorphism (g is an isomorphism except at gadget $Z(u)$, and Spoiler would need more than one pebble there to uncover the difference). The move of Spoiler induces an update for the i th cop in the cops and robber game, which yields a response of the robber according to his winning strategy, i.e. a move along a cop-free path P to some vertex v . Hence, as shown above, the bijection $g' := g \circ \eta_P$ respects all

pebbled pairs of elements and is good except at v . Since $Z(v)$ is cop-free (and hence not pebbled), the claim follows. Q.E.D.

Theorem 7.17. $\text{FPC} \subsetneq \text{PTIME}$ on every class of graphs which contains CFI-graphs $X(G)$ and $\tilde{X}(G)$ for graphs G of arbitrary large treewidth.

In fact, Grohe and Marino proved that $\text{FPC} \equiv \text{PTIME}$ on every class of graphs with bounded treewidth. Their theorem allows us to reformulate the result in a very neat way.

We first observe that the treewidth of $X(G)$ is bounded by $(O)(\text{tw}(G))$: from a tree-decomposition of G one obtains a tree decomposition of $X(G)$ by replacing in all bags the vertices by their corresponding gadgets. Furthermore, the size of a gadget $Z(v)$ in $X(G)$ is bounded by $(4\Delta(G) \cdot 2^{\Delta(G)-1}) \in \mathcal{O}(\Delta(G))$. Now let G_n be the $n \times n$ grid, then $\text{tw}(G_n) = n$, $\Delta(G_n) = 4$ and

$$\text{tw}(X(G)) \leq (4\Delta(G) \cdot 2^{\Delta(G)-1}) \text{tw}(G_n) = 24n \in \mathcal{O}(|G|).$$

For a function $f : \mathbb{N} \rightarrow \mathbb{N}$ we define the class of graphs

$$\text{TW}_f := \{G : \text{tw}(G) \leq f(|G|)\}.$$

Theorem 7.18. $\text{FPC} \equiv \text{PTIME}$ on TW_f if, and only if, $f \in \mathcal{O}(1)$.

Proof. The direction from right to left is mentioned theorem due to Marino and Grohe. For the other direction, assume $f \notin \mathcal{O}(1)$; then for every $n > 0$, there exists $k > |X(G_n)|$ with $f(k) \geq 24n$. Hence, TW_f contains $X(G_n)$ and $\tilde{X}(G_n)$ for every $n \geq 0$. Q.E.D.