4. Fixed-Point Logic with Counting

This chapter is devoted to the introduction and analysis of the natural extensions of the fixed-point logics FP and PFP that have expressive means for cardinality properties.

• The actual formalization of fixed-point logics with counting, FP+C and PFP+C, in a two-sorted framework is given in Section 4.1.

• In Section 4.2 the relation of FP+C and PFP+C with the $C^k_{\infty\omega}$ and with the C^k-invariants is investigated. In particular we obtain the analogue of the first theorem of Abiteboul and Vianu (Theorem 3.22 above) in the presence of counting. In contrast with the second theorem of Abiteboul and Vianu (Theorem 3.24) we here find that FP+C *is* the polynomial restriction of PFP+C.

• Section 4.3 deals with the separation result FP+C $\subsetneq$ PTIME, which is due to Cai, Fürer and Immerman, in a framework that lends itself to relativization. In restriction to classes with certain closure properties FP+C can only capture PTIME if some I_{C^k} provides a complete invariant *up to isomorphism* (equivalently, if some $C^k_{\infty\omega}$ coincides with $L_{\infty\omega}$) over this class.

• Section 4.4 summarizes some results on equivalent characterizations of the expressive levels of FP+C and PFP+C.

As pointed out in the introduction, first-order logic at first sight suffers from *two* independent shortcomings over finite structures: it completely lacks mechanisms to model recursion — the fixed-point operations provided in FP and PFP answer this requirement; and it also lacks expressive means to assess cardinalities of definable sets. The latter defect is obviously overcome automatically together with the former over ordered structures. By the theorems of Immerman, Vardi and Abiteboul, Vardi, Vianu, FP and PFP capture PTIME and PSPACE over ordered structures. In particular all PTIME, respectively PSPACE, properties of cardinalities are expressible in FP, respectively PFP, over ordered structures. Not so in the case of not necessarily ordered structures: in fact the most obvious examples that FP and PFP do not correspond to standard complexity classes in the general case all involve counting. Over pure sets for instance FP, PFP and even $L^\omega_{\infty\omega}$ collapse to first-order logic and cannot express low complexity cardinality properties like "there is an even number of elements". For some time therefore it had been conjectured, mainly by Immerman, that FP enriched with counting might capture PTIME in the general case. This expectation was later disproved by Cai, Fürer and Immerman, who showed that not even $C^{\omega}_{\infty\omega}$ comprises all of PTIME. There remains good motivation to study the extensions of FP and PFP by expressive means for counting, however.

(a) FP and PFP are successful extensions of first-order logic and capture an interesting notion of *relational recursion* on finite structures even in the absence of order. Without order, however, they do not add to the power of first-order with respect to cardinality properties. Many natural structural properties involve counting in addition to relational recursion.
(b) It is reasonable to treat the two obvious defects of first-order logic on an equal footing and to investigate natural levels of expressiveness that address both defects.
(c) As indicated in the previous chapter, the relationship between FP and PFP on the one hand and $L^{\omega}_{\infty\omega}$ on the other leads to valuable insights into the nature of relational recursion on finite structures. FP and PFP *are*, in some intuitive sense, PTIME *and* PSPACE *in the world of* $L^{\omega}_{\infty\omega}$. $C^{\omega}_{\infty\omega}$ is a natural richer and still well-behaved fragment of $L_{\infty\omega}$. In particular $C^{\omega}_{\infty\omega}$ shares with $L^{\omega}_{\infty\omega}$ the benefit of elegant game characterizations and the existence of PTIME computable invariants. It is natural therefore to expect appropriate counting extensions of FP and PFP to *represent* PTIME *and* PSPACE *in the world of* $C^{\omega}_{\infty\omega}$.
(d) It turns out that fixed-point logics with counting represent robust levels of expressive power in the sense that the semantic strength proves to be independent of several choices in the actual formalization. More importantly they offer a number of interesting equivalent characterizations.
(e) Finally we shall see that some properties of the counting extensions resemble those found for FP and PFP themselves only in the ordered case. Roughly speaking, with counting one is closer to the ordered case.

4.1 Definition of FP+C and PFP+C

The natural modelling for the counting extensions uses two-sorted structures. The given relational structure forms the first sort, an *ordered numerical domain* the second. In this way counting terms that take values in the numerical domain can naturally be introduced. The two-sorted structures can of course in the standard way be encoded in a one-sorted framework with extra unary predicates to denote the different universes. We shall at some points appeal to this possibility. For the basic formalizations, however, the two-sorted picture is easier to handle and intuitively neater. Let τ be finite and relational as usual.

Definition 4.1. *Let * be the functor that takes $\mathfrak{A} \in \text{fin}[\tau]$ to the two-sorted structure $\mathfrak{A}^*$ which is the disjoint union of $\mathfrak{A}$ itself for the first sort and the canonical ordered structure of size $|A| + 1$ for the second sort.*

$$\mathfrak{A}^* := \mathfrak{A} \,\dot{\cup}\, (n+1, <^{n+1}) \quad \textit{where } n = |A|.$$

Recall that we identify $n+1$ with $\{0, \ldots, n\}$. Let $\text{fin}[\tau]^* := \{\mathfrak{A}^* \mid \mathfrak{A} \in \text{fin}[\tau]\}$.

We apply the following formalism to the two-sorted structures in $\text{fin}[\tau]^*$. Variable symbols x, y, z, ... range over the elements of the first sort, variables ν, μ, ... range over the second sort. Of a second-order variable X we say that it is of type (r_1, r_2) if it ranges over subsets of $(\text{first sort})^{r_1} \times (\text{second sort})^{r_2}$. All second-order variables come with a definite typing in this sense.

We consider first-order logic and its extensions by FP- and PFP-operators over $\text{fin}[\tau]^*$. The first-order constructors comprise

- the formation of atomic expressions, which have to respect the type of second-order variables in the obvious way,
- boolean connectives and
- quantifications with respect to first-order variables of each type.

For the fixed-point operators we admit the most general kind of fixed-point generations in the two-sorted framework by allowing fixed-point variables X of arbitrary mixed types. Otherwise no changes are necessary to accommodate fixed-point operations over the $\mathfrak{A}^*$. Compare Section 1.3.3 and Definitions 1.22 and 1.23. Let for instance $\varphi(X, \overline{x}, \overline{\nu})$ be in the indicated free variables, where X is of type (r_1, r_2) and $\overline{x}$ and $\overline{\nu}$ are tuples of r_1, respectively r_2, distinct variables for elements of the first, respectively second, sort.

Over each $\mathfrak{A}^* \in \text{fin}[\tau]^*$ the formula φ induces the following mapping $F_\varphi^{\mathfrak{A}^*}$ ($\mathcal{P}$ denotes the power set).

$$\begin{array}{rcl} F_\varphi^{\mathfrak{A}^*} \colon \mathcal{P}\Big(A^{r_1} \times \{0, \ldots, |A|\}^{r_2}\Big) & \longrightarrow & \mathcal{P}\Big(A^{r_1} \times \{0, \ldots, |A|\}^{r_2}\Big) \\ P & \longmapsto & \Big\{(\overline{a}, \overline{m}) \;\Big|\; \mathfrak{A}^* \models \varphi[P, \overline{a}, \overline{m}]\Big\}. \end{array}$$

The semantics of formulae $[\text{PFP}_{X,\overline{x}\overline{\nu}}\varphi(X, \overline{x}, \overline{\nu})]\overline{x}\,\overline{\nu}$ and $[\text{FP}_{X,\overline{x}\overline{\nu}}\varphi(X, \overline{x}, \overline{\nu})]\overline{x}\,\overline{\nu}$ is defined in terms of the partial, respectively inductive or inflationary, fixed points of F_φ just as in the one-sorted case.

Definition 4.2. *Let $L^*_{\omega\omega}$ be two-sorted first-order logic for *-structures. Similarly $L^{\omega\,*}_{\infty\omega}$ is that fragment of infinitary logic for two-sorted *-structures that consists of formulae using only finitely many first-order variables (of either sort).* FP* *and* PFP* *stand for the two-sorted variants of fixed-point and partial fixed-point logic for these two-sorted structures.*

Note that these logics admit formulae with free first-order variables of both sorts, or, where applicable, also free second-order variables of mixed type. We ultimately only consider formulae that are free over first-order variables of the first sort and define global relations over the original relational structures. This might be regarded as the *standard part* of the semantics for these logics. All considerations about the expressive power of these logics concern these standard parts. A statement like $C^{\omega}_{\infty\omega} \supseteq \mathrm{FP}^*$, for instance, means that any global relation over fin[τ] that is FP*-definable as a global relation over the first sort is $C^{\omega}_{\infty\omega}$-definable. Formulae with other free variables are important, however, for the inductive generation of formulae and accordingly play some rôle in particular in syntactic arguments by induction.

A technical comment is in order with respect to the standard one-sorted modelling of two-sorted structures. In the sequel we shall want to apply results that formally deal with one-sorted structures also in the present two-sorted formalization. Rather than reproving them in a tedious adaptation of the standard arguments one may directly apply them on the basis of the following remark.

Remark 4.3. *For* $\mathcal{L} = L_{\omega\omega}$, $L^{\omega}_{\infty\omega}$, FP, PFP *and with* $\mathcal{L}^* = L^*_{\omega\omega}$, $L^{\omega\,*}_{\infty\omega}$, FP*, PFP* *according to Definition 4.2: the expressive power of* $\mathcal{L}$ *on the standard one-sorted encodings of structures in* fin$[\tau]^*$ *is the same as that of* $\mathcal{L}^*$.

Sketch of Proof. The argument is via mutual simulations between the one-sorted and the two-sorted frameworks.
i) First-order constructors. Consider first the simulation of the two-sorted framework in the one-sorted encodings, where the i-th sort is described by a unary predicate U_i. The distinction between first-order variables of different sorts is faithfully simulated through relativizations to the respective subdomains. Conversely, a formula $\varphi(x_1, \dots, x_r)$ of the one-sorted framework, whose first-order variables range over the combined domain $U_1 \dot{\cup} U_2$, translates into a tuple of 2^r formulae φ_s, $s \subseteq \{1, \dots, r\}$, of the two-sorted framework — one for each possible typing. For instance if $s = \{1, 2\}$, then $\varphi_s = \varphi_s(x_1, x_2, \nu_3, \dots, \nu_r)$ takes care of the case that just x_1 and x_2 get interpreted over U_1. The inductive definition of the φ_s is straightforward. For instance, if $\psi(x_1, \dots, x_{r-1}) = \exists x_r \varphi(x_1, \dots, x_r)$, then $\psi_{\{1,2\}} = \exists x_r \varphi_{\{1,2,r\}} \vee \exists \nu_r \varphi_{\{1,2\}}$.
ii) Second-order variables and fixed-point processes. A second-order variable X that is of type (r_1, r_2) over the two-sorted structures is simulated over their one-sorted encodings by a second-order variable of arity $r_1 + r_2$ which can easily be relativized to interpretations of the correct type. Fixed-point processes carry over directly. In the other direction consider an r-ary second-order variable X over the one-sorted encodings. Since its interpretations do not come with a fixed typing, it has to be modelled in general by a tuple of 2^r second-order variables $(X_s)_{s \subseteq \{1,\dots,r\}}$, one for each possible typing. We think of the original X as the union of the X_s where, for instance, $X_{\{1,2\}}$ is the collection of tuples in X whose first two components come from U_1. Obviously X and the X_s are first-order interdefinable (over the one-sorted encodings). A fixed-point process involving X naturally translates into a simultaneous fixed-point process for a system of formulae. In this system there is one formula $\varphi_s\big((X_t)_{t \subseteq \{1,\dots,r\}}\big)$ in first-order variables typed according to s,

for each s. The resulting fixed points of systems can be recast into ordinary fixed points using standard techniques as discussed in Example 1.27. $\square$

This observation also implies that the usual semantic inclusions carry over to the two-sorted framework.

Remark 4.4. $L^*_{\omega\omega} \subseteq \mathrm{FP}^* \subseteq \mathrm{PFP}^* \subseteq L^{\omega\,*}_{\infty\omega}$.

The functor $^*: \mathrm{fin}[\tau] \longrightarrow \mathrm{fin}[\tau]^*$ is isomorphism preserving: $\mathfrak{A}^* \simeq \mathfrak{A}'^*$ if and only if $\mathfrak{A} \simeq \mathfrak{A}'$. Similarly it preserves the substructure relation. It does not, however, preserve definability of substructures even at the atomic level. As a consequence, FP* and PFP* do not have the relativization property. For a simple example consider evenness. Evenness of the universe is obviously definable in FP*: $|A|$ is even if the ordered second sort of $\mathfrak{A}^*$ has an odd number of elements, and FP*-recursion over the second sort suffices for checking this. Evenness of a unary predicate $U \in \tau$, however, is not in FP*. The straightforward adaptation of the standard game argument shows that evenness of $U \subseteq A$ is not even definable in $L^{\omega\,*}_{\infty\omega}$. In a sense, only the cardinality of the universe has yet been made available in the ordered numerical sort. To introduce counting and to remedy the defects just pointed out, it suffices to render the cardinalities of definable subsets over the first sort definable over the second sort. We present below two equivalent ways of doing so. The first approach introduces *counting terms* in a straightforward way. The other one — more elegant maybe from a model theoretic view — uses the extension by the Härtig quantifier.

Counting terms. Counting terms link the two sorts so that the second, numerical sort can be used for talking about the size of definable subsets. It suffices to consider unary subsets of the first sort, for reasons discussed below.

Definition 4.5. *With each formula φ and any variable x of the first sort associate a counting term:*

$$t := \#_x\varphi(x)$$

of the second sort. Put $\mathrm{free}(t) = \mathrm{free}(\varphi) \setminus \{x\}$. *If $(\mathfrak{A}^*, \Gamma)$ interprets all free variables of φ apart from x, then the interpretation of t in $(\mathfrak{A}^*, \Gamma)$ is that element of the second sort that describes the size of the predicate $\varphi[\mathfrak{A}^*, \Gamma]$ defined by φ:*

$$t^{\mathfrak{A}^*,\Gamma} := \Big|\big\{a \in A \mid (\mathfrak{A}^*, \Gamma) \models \varphi[a]\big\}\Big|.$$

To obtain fixed-point logic with counting, we simultaneously close first-order logic $L^*_{\omega\omega}$ for the two-sorted structures under the FP*-constructor, the formation of counting terms and substitution of these for variables of the second sort. The formal definition of the syntax would be via a combined inductive generation of formulae and terms.

Definition 4.6. *Let* FP+C *be the smallest extension of* FP* *that is closed under formation and substitution of counting terms and the* FP*-*constructor.* PFP+C *is the corresponding closure with respect to* PFP*.

Clearly FP+C $\subseteq$ PTIME and PFP+C $\subseteq$ PSPACE.

Using the Härtig quantifier instead. The Härtig quantifier, cf. Definition 1.53, expresses cardinality equality. Its semantics extends naturally to two-sorted structures. Over the $\mathfrak{A}^*$ it may be used to define counting-terms:

$$\begin{array}{ll} \nu = \#_x\varphi(x) & \text{is equivalent with} \\ Q_{\mathrm{H}}\Big((x;\varphi);(\mu;\psi)\Big) & \text{for } \psi(\mu,\nu) = \mu < \nu. \end{array}$$

Denote by $\mathrm{FP}(Q_{\mathrm{H}})^*$ and $\mathrm{PFP}(Q_{\mathrm{H}})^*$ the logics that result from adjoining the Härtig quantifier in the two-sorted framework. It turns out that these provide equivalent characterizations for FP+C and PFP+C, respectively.

Lemma 4.7. PFP+C $\equiv$ $\mathrm{PFP}(Q_{\mathrm{H}})^*$ *and* FP+C $\equiv$ $\mathrm{FP}(Q_{\mathrm{H}})^*$.

Proof. We point out that the statement of Remark 4.3 extends to the extensions of FP and PFP by the Härtig quantifier. Adjoining the Härtig quantifier over the one-sorted encodings of two-sorted structures, we formally gain cardinality equalities for mixed-sorted unary predicates. These however are dissolved into equalities for the sums of cardinalities for two unary pure-sorted predicates each. Sums over the second sort, however, are definable in FP over the ordered second sort since they are PTIME computable.

For the proof of the lemma note that the inclusions "$\subseteq$" follow directly from the definability of counting terms through the Härtig quantifier. Consider then the converse inclusion for FP. An application of the Härtig quantifier may involve two predicates over the first-sort — this case translates into an equality for the corresponding counting terms directly. It may also involve at least one predicate over the second sort — but over ordered domains, values of counting terms of type $\#_\nu\varphi$ are even FP-definable since they are PTIME computable. □

FP+C and PFP+C turn out to be very robust with respect to the formal details concerning the introduction of counting terms. For example, it is natural to allow counting not only for unary predicates but also in higher arities and over mixed sorts. We have just seen that unary counting over the second sort is for free. The reason for this robustness is that in FP+C we already have the full power of PTIME operations over the second sort. This is at the root of the following model theoretic statement of robustness. It should be noted that a corresponding counting extension of first-order logic does not at all share these properties, see Example 4.13 below. For the notion of generalized interpretations and closure with respect to these compare Definitions 1.44 and 1.48 in Section 1.5.

Proposition 4.8. FP+C *and* PFP+C *are closed with respect to generalized interpretations.*

Proof. Consider for instance FP+C. The statement to be proved is the following. Let $\mathfrak{i}$ be some FP+C-definable generalized (σ, τ)-interpretation, functorially $\mathfrak{i}\colon \mathrm{fin}[\tau] \to \mathrm{fin}[\sigma]$. Let R be some FP+C-definable global relation over $\mathrm{fin}[\sigma]$. Then the global relation $\mathfrak{i}(R)$ over $\mathrm{fin}[\tau]$ whose value over $\mathfrak{A}$ is the interpretation over $\mathfrak{A}$ of $R^{\mathfrak{i}(\mathfrak{A})}$ has to be FP+C-definable as well. Since we know that FP and PFP have the required closure properties, it suffices to prove the following.

(∗) Each definable interpretation of σ-structures over $\mathrm{fin}[\tau]$ induces a definable interpretation of the corresponding two-sorted structures in $\mathrm{fin}[\sigma]^*$ over $\mathrm{fin}[\tau]^*$. This interpretation is such that counting terms for the interpreted $\mathrm{fin}[\sigma]^*$-structures are FP+C-definable over the parent structures in $\mathrm{fin}[\tau]^*$.

Sufficiently large numerical domains are interpretable in powers of the given numerical domain. The set of s-tuples over $n+1$ together with the first-order definable lexicographic ordering provides an interpretation of $\big((n+1)^s, <\big)$ over $(n+1, <)$ as always $(n+1)^s \geqslant n^s + 1$. This numerical domain is sufficiently large to provide the second sort for interpretations over the s-th power. The numerical value represented by an s-tuple $\overline{m}$ in $(n+1)^s$ is the number of lexicographic predecessors of $\overline{m}$: $\big|\{\overline{m}' | \overline{m}' <_{\mathrm{lex}} \overline{m}\}\big|$. Having these numerical domains, (∗) reduces to the following lemma: FP+C suffices to simulate counting terms over interpretations in powers and quotients. □

Lemma 4.9. *The analogues of counting terms for counting in higher arity and for counting modulo definable congruences (counting equivalence classes) are definable over* $\mathrm{fin}[\tau]^*$ *in* FP+C.

Proof. The claim for higher arity counting means that for $\varphi(x_1, \ldots, x_s)$ in FP+C (where other variables are suppressed without loss of generality) there is a formula $\psi(\nu_1, \ldots, \nu_s)$ in FP+C such that

$$\mathfrak{A}^* \models \psi[\overline{m}] \quad \Longleftrightarrow \quad \big|\varphi[\mathfrak{A}^*]\big| = \big|\{\overline{m}' | \overline{m}' <_{\mathrm{lex}} \overline{m}\}\big|.$$

Consider for instance the binary case, a formula $\varphi(x, y)$. For each m, the number of x such that there are exactly m many y satisfying φ with that x is $t(\mu) = \#_x\big(\#_y \varphi(x,y) = \mu\big)$, where μ is the second-sort variable for m. But obviously the desired lexicographic representation of the number $l = \big|\{(x,y)|\varphi\}\big|$ is PTIME computable in terms of the function $m \mapsto t(m)$ through

$$l = \sum_m m\, t(m).$$

The graph of the function $m \mapsto t(m)$ is FP+C-definable over the second sort so that FP+C-definability of l follows immediately.

Counting with respect to a definable congruence, or the lexicographic representation of the number $l = \left|\varphi[\mathfrak{A}^*]/\psi[\mathfrak{A}^*]\right|$, is treated analogously. Without loss of generality let now $\varphi = \varphi(x)$ be unary, $\psi = \psi(x,y)$ binary. Here l is PTIME computable from the function $m \mapsto t(m)$ where $t(\mu) = \#_x(\varphi(x) \wedge \#_y(\varphi(y) \wedge \psi(x,y)) = \mu)$, such that $t(m)$ is the number of elements whose ψ-class in φ has exactly m elements and

$$l = \sum_m m^{-1} t(m).$$

□

Example 4.10. All (even quotient) cardinality Lindström quantifiers (see Definitions 1.52 and 1.54) that are based on PTIME computable numerical predicates are expressible in FP+C. This is an obvious consequence of the above fact that FP+C has definable counting terms for counting in arbitrary arities and with respect to definable congruences together with the Immerman-Vardi theorem applied to fixed-point definability over the second sort.

Example 4.11. Since in particular the Rescher quantifier (Definition 1.53) is definable in FP+C we obtain from Lemma 2.22 that the stable colouring of graphs is FP+C-definable. It similarly follows from Proposition 3.6 that the relational parts of the the C^k-invariants are FP+C-interpretable as quotients over the k-th power. This is further explored in the next section.

Example 4.12. $\equiv^{C^k}$ is in FP+C, just as $\equiv^{L^k}$ is in FP according to Corollary 3.15. This is easier to see than the stronger claim made in Proposition 3.10 about definability in FP(Q_{H}), since one may here argue directly with interpretability of the relational parts of the I_{C^k} together with availability of counting terms to check equality for the weights.

Aside on first-order logic with unary counting. As pointed out above, first-order logic is far more sensitive to slight changes in the definition of a "counting extension" than FP and PFP are. This is not surprising since the robustness of FP+C and PFP+C is due to their recursive power over the second sort. Let for the considerations of the following example *first-order logic with unary counting* be defined as the closure of $L^*_{\omega\omega}$ with respect to the formation and substitution of counting terms in the sense of Definition 4.5.

Example 4.13. First-order logic with unary counting does not capture binary counting. Consider $\tau = \{U_1, U_2, U_3\}$ consisting of three unary predicates. Let Q be the class of those τ-structures whose universe is partitioned into three disjoint sets by the U_i. Let always m_i stand for the cardinality of U_i, and $n = m_1 + m_2 + m_3$ for the overall size of $\mathfrak{A} \in Q$. The tuple (m_1, m_2, m_3) characterizes $\mathfrak{A}$ up to isomorphism, of course. Let $Q_0 \subseteq Q$ be the subclass defined by the condition $m_2 = m_1^2$. Clearly Q_0 is definable in first-order logic with counting terms for binary predicates: one need merely equate the cardinalities of the first-order definable predicates $\{(x,y)|x = y \wedge U_2x\}$ and $\{(x,y)|U_1x \wedge U_1y\}$.

We claim that Q_0 is not definable in $L^*_{\omega\omega}$ with unary counting terms. Call this logic $\mathcal{L}$ for the purposes of this proof. The proof involves a reduction of definability in $\mathcal{L}$ to ordinary first-order definability over the second, arithmetical sort of the $\mathfrak{A}^*$ expanded with just a fixed finite number of constants for some particular values of counting terms. Standard Ehrenfeucht-Fraïssé arguments for linear orderings then apply to show that Q_0 cannot be separated from $Q \setminus Q_0$ by these first-order means.

A trivial automorphism argument will be used repeatedly. If $\overline{a}$ and $\overline{a}'$ are such that $\mathrm{atp}_{\mathfrak{A}}(\overline{a}) = \mathrm{atp}_{\mathfrak{A}}(\overline{a}')$ then there is an automorphism of $\mathfrak{A}^*$ which maps $\overline{a}$ to $\overline{a}'$ and fixes the second sort of $\mathfrak{A}^*$ pointwise. It follows that

(i) for $\varphi(\overline{x},\overline{\nu}) \in \mathcal{L}$ and fixed interpretation $\overline{m}$ for $\overline{\nu}$ over $\mathfrak{A}^*$, the predicate $\varphi[\mathfrak{A}^*,\overline{m}] = \{\overline{a}\,|\,\mathfrak{A}^* \models \varphi[\overline{a},\overline{m}]\}$ is a union of sets $\theta[\mathfrak{A}]$ for $\theta \in \mathrm{Atp}(\mathfrak{A};k)$, ($k$ the arity of $\overline{x}$).
(ii) for $\theta \in \mathrm{Atp}(\mathfrak{A};k)$ and $\Theta \subseteq \mathrm{Atp}(\mathfrak{A};k)$ the counting values

$$t(\theta,\Theta)^{\mathfrak{A}} = \left|\{b \in A \mid \mathrm{atp}\left(\overline{a}\tfrac{b}{1}\right) \in \Theta\}\right|$$

for $\overline{a} \in \theta[\mathfrak{A}]$ only depend on $\mathfrak{A}$, θ and Θ (and not on $\overline{a} \in \theta[\mathfrak{A}]$).
(iii) for θ, Θ as above and for all $\mathfrak{A} \in Q$ with sufficiently large $m_i = |U_i^{\mathfrak{A}}|$, $t(\theta,\Theta)$ is of the form $\sum_{i\in s} m_i \pm d$, where $s \subseteq \{1,2,3\}$ and $0 \leqslant d \leqslant k$, s and d depending only on θ and Θ.

Consider the second sort of $\mathfrak{A}^*$, for $\mathfrak{A} \in Q$, as equipped with parameters $\overline{t}$ for the values of all $t(\theta,\Theta)^{\mathfrak{A}}$ (for fixed k, as appropriate).

Claim. For each $\varphi(\overline{x},\overline{\nu}) \in \mathcal{L}$ and $\theta \in \mathrm{Atp}(\mathfrak{A};k)$ there is a $<$-formula $\underline{\varphi}_\theta(\overline{\nu},\overline{\mu})$ in first-order logic (for the the second sort) such that for all $\mathfrak{A} \in Q$ with sufficiently large m_i and for all interpretations $\overline{m}$ for the $\overline{\nu}$:

$$\theta[\mathfrak{A}] \subseteq \varphi[\mathfrak{A}^*,\overline{m}] \quad \Leftrightarrow \quad \mathfrak{A}^* \models \underline{\varphi}_\theta[\overline{m},\overline{t}].$$

This claim is justified inductively. The atomic cases and boolean connectives are trivially dealt with.

If $\varphi = \exists x_j \psi(\overline{x},\overline{\nu})$, then $\underline{\varphi}_\theta$ is the disjunction over all $\underline{\psi}_{\theta'}$ with $\theta' \in \mathrm{Atp}(\tau;k)$ such that θ' and θ agree on $\{x_1,\dots,x_k\} \setminus \{x_j\}$. For $\varphi = \exists\nu\psi(\overline{x},\overline{\nu})$ one can simply take $\underline{\varphi}_\theta = \exists\nu\underline{\psi}_\theta$.

Finally let $\varphi = \#_{x_1}\psi(\overline{x},\overline{\nu}) = \nu$. Then $\theta[\mathfrak{A}] \subseteq \varphi[\mathfrak{A}^*,\overline{m},m]$ if $t(\theta,\Theta)^{\mathfrak{A}} = m$ for $\Theta = \{\theta' \in \mathrm{Atp}(\tau;k) \mid \theta'[\mathfrak{A}] \subseteq \psi[\mathfrak{A}^*,\overline{m}]\}$. But by the inductive hypothesis $\Theta = \{\theta' \mid \mathfrak{A}^* \models \underline{\psi}_{\theta'}[\overline{m},\overline{t}]\}$. The equation $t(\theta,\Theta) = m$ can be put into the desired form through a distinction of cases: $t(\theta,\Theta) = \nu$ is equivalent with the disjunction of the following formulae, over all subsets $\Theta' \subseteq \mathrm{Atp}(\tau;k)$:

$$\bigwedge_{\theta'\in\Theta'} \underline{\psi}_{\theta'}[\overline{\nu},\overline{\mu}] \;\wedge \bigwedge_{\theta'\notin\Theta'} \neg\underline{\psi}_{\theta'}[\overline{\nu},\overline{\mu}] \;\wedge\; t(\theta,\Theta') = \nu\,.$$

This proves the claim.

For sentences $\varphi \in \mathcal{L}$ it follows that there is a formula $\underline{\varphi}$ of $L_{\omega\omega}[<]$ such that for all $\mathfrak{A} \in Q$: $\mathfrak{A}^* \models \varphi \Leftrightarrow (\{0, \dots, |A|\}, <, \overline{t}^{\mathfrak{A}}) \models \underline{\varphi}$.

The standard Ehrenfeucht-Fraïssé analysis of linear orderings shows that no first-order formula of quantifier rank q can distinguish

$$(n+1, <, t_1, \dots, t_l) \text{ from } (n+1, <, t'_1, \dots, t'_l)$$

$$\begin{aligned} \text{if } \ & 0 = t_1 < t_2 < \dots < t_l = n, \\ & 0 = t'_1 < t'_2 < \dots < t'_l = n, \end{aligned}$$

$$\text{and if for all } i, j\text{: either } |t_j - t_i| = |t'_j - t'_i| \text{ or } |t_j - t_i|, |t'_j - t'_i| \geqslant 2^q.$$

By (iii) above we see that this degree of similarity is achieved for structures $(n+1, <, \overline{t}^{\mathfrak{A}})$ and $(n+1, <, \overline{t}^{\mathfrak{A}'})$ whenever 0, m_1, m_2, m_3, and $n+1 = m_1 + m_2 + m_3 + 1$ are spaced sufficiently far apart. Therefore, no first-order formula can separate those $(n+1, <, \overline{t}^{\mathfrak{A}})$ for $\mathfrak{A} \in Q_0$ from those for $\mathfrak{A} \notin Q_0$, and Q_0 cannot be definable in $\mathcal{L}$ either.

4.2 FP+C and the C^k-Invariants

We saw in Section 3.4.1 that interpretability of the L^k-invariants in fixed-point logic on the one hand and representability of fixed-point processes over the invariants on the other hand lead to characterizations of the expressive power of FP and PFP in terms the I_{L^k}. An important aspect of this characterization is the reduction to ordered domains. FP and PFP over not necessarily ordered structures can be analyzed in terms of FP and PFP over the linearly ordered invariants. This section is devoted to the corresponding analysis for fixed-point logics with counting.

The first lemma concerns FP+C-interpretability of the I_{C^k}-invariants. Essentially this is a restatement of the definability properties of the I_{C^k} expressed in Proposition 3.6 above — now put in terms of FP+C.

Lemma 4.14. *$I_{C^k}(\mathfrak{A})$ is* FP+C*-interpretable over $\mathfrak{A}^*$. More precisely all the following are* FP+C*-interpretable:*

(i) the relational part of I_{C^k} as a quotient over the k-th power over the first sort.

(ii) I_{C^k} as a whole (and being a standard structure) over the second sort.

(iii) the natural projection from the quotient interpretation of the relational part of I_{C^k} over the first sort to its representation over the second sort.

Proof. Proposition 3.6 applies to show (i) since $\mathrm{FP}(Q_{\mathrm{R}}) \subseteq$ FP+C. FP* itself suffices to define the natural projection from the pre-ordering in this interpretation of the relational part of I_{C^k} to the ordered quotient structure over the second sort (iii). Definability of the weight functions through simple counting terms as stated in Proposition 3.6 completes the interpretability of the full invariant as expressed in (ii). □

Let $I^*_{C^k}(\mathfrak{A}) := I_{C^k}(\mathfrak{A}^*)$ stand for the C^k-invariant of $\mathfrak{A}^* \in \text{fin}[\tau]^*$, more precisely of the standard one-sorted encoding of $\mathfrak{A}^*$.

Lemma 4.15. *The $I^*_{C^k}$ are* FP*-interpretable over the I_{C^k}.*

Proof. Since we are dealing with ordered structures it suffices to show that there is a PTIME algorithm that computes $I^*_{C^k}(\mathfrak{A})$ from $I_{C^k}(\mathfrak{A})$. But the inductive generation of $I^*_{C^k}(\mathfrak{A})$ is obviously in PTIME and requires no other data than those encoded in $I_{C^k}(\mathfrak{A})$. The initial stage for instance is based on some fixed ordering of the atomic types of k-tuples in $\mathfrak{A}^*$. Since $\mathfrak{A}^*$ is the disjoint union of $\mathfrak{A}$ with the linear ordering $(|A| + 1, <)$, these atomic types can be presented by pairs of atomic types, one in vocabulary τ for the components in the first sort and one in vocabulary $<$ for the components in the second sort. Note that all the relevant information about r-tuples over $\mathfrak{A}$ for $r < k$ is also encoded in $I_{C^k}(\mathfrak{A})$ since the C^k-type of a tuple $(x_1, \ldots, x_r)$ is encoded by the C^k-type of the k-tuple $(x_1, \ldots, x_1, x_1, \ldots, x_r)$ with $r - k$ additional entries x_1. In this fashion the inductive steps in the generation of $I^*_{C^k}(\mathfrak{A})$ are easily simulated over $I_{C^k}(\mathfrak{A})$. □

The lemma is in fact a special case of the following more general observation that can be proved along the same lines. The statement admits further generalizations in the style of Feferman-Vaught Theorems for the L^k- and C^k-theories of finite structures.

Remark 4.16. *The L^k- and C^k-invariants are modular with respect to disjoint unions and direct products in the sense that, for example for I_{C^k} and for disjoint unions, there is a* PTIME *function Σ such that for all $\mathfrak{A}, \mathfrak{B} \in \text{fin}[\tau]$:*

$$I_{C^k}\big(\mathfrak{A} \,\dot{\cup}\, \mathfrak{B}\big) = \Sigma\Big(I_{C^k}(\mathfrak{A}), I_{C^k}(\mathfrak{B})\Big).$$

This implies also that $I_{C^k}\big(\mathfrak{A} \,\dot{\cup}\, \mathfrak{B}\big)$ is FP*-interpretable over the disjoint union of $I_{C^k}(\mathfrak{A})$ and $I_{C^k}(\mathfrak{B})$.*

PTIME$\big(I_{C^k}\big)$ and PSPACE$\big(I_{C^k}\big)$ are defined in analogy with Definition 3.20:

Definition 4.17. PTIME$\big(I_{C^k}\big)$ *and* PSPACE$\big(I_{C^k}\big)$ *stand for the classes of all those queries that are* PTIME, *respectively* PSPACE, *computable in terms of the I_{C^k}.*

More precisely, a boolean query Q on fin$[\tau]$ *is in* PTIME$\big(I_{C^k}\big)$ *if membership of $\mathfrak{A}$ in Q is a* PTIME *property of $I_{C^k}(\mathfrak{A})$. A similar characterization can be applied to global relations (of arity at most k) using the extensions of the invariants to the* fin$[\tau; r]$.

As in the corresponding treatment of the I_{L^k} a query is in PTIME$\big(I_{C^k}\big)$ respectively PSPACE$\big(I_{C^k}\big)$ if it is $C^k_{\infty\omega}$-definable and its natural representation over the relational part of the I_{C^k} can be computed in PTIME, respectively PSPACE over the I_{C^k}. Logically these classes can further be identified with classes $\mathrm{FP}(I_{C^k})$ and $\mathrm{PFP}(I_{C^k})$ since FP and PFP capture PTIME and PSPACE over the ordered I_{C^k}:

$$\begin{aligned}\mathrm{FP}(I_{C^k}) &\equiv \mathrm{PTIME}(I_{C^k}),\\ \mathrm{PFP}(I_{C^k}) &\equiv \mathrm{PSPACE}(I_{C^k}).\end{aligned}$$

Syntactically the formulae of $\mathrm{FP}(I_{C^k})$ or $\mathrm{PFP}(I_{C^k})$ are FP*-formulae, respectively PFP*-formulae, in terms of the interpreted I_{C^k}. These logics may thus be regarded as fragments of FP+C or PFP+C. See the proof of the following theorem.

Theorem 4.18. *With the* $\mathrm{FP}(I_{C^k})$ *and* $\mathrm{PFP}(I_{C^k})$ *as characterized:*

$$\begin{aligned}\mathrm{FP{+}C} &\equiv \textstyle\bigcup_k \mathrm{FP}(I_{C^k}) &&\equiv \textstyle\bigcup_k \mathrm{PTIME}(I_{C^k})\\ \mathrm{PFP{+}C} &\equiv \textstyle\bigcup_k \mathrm{PFP}(I_{C^k}) &&\equiv \textstyle\bigcup_k \mathrm{PSPACE}(I_{C^k}).\end{aligned}$$

Proof. We prove the equivalences between the logical characterizations. The arguments for FP+C and PFP+C are completely analogous. Consider FP+C. By Lemma 3.19 $\mathrm{FP}(Q_{\mathrm{H}}) \subseteq \bigcup_k \mathrm{FP}(I_{C^k})$. An application to the one-sorted encodings of structures in $\mathrm{fin}[\tau]^*$ yields

$$\mathrm{FP}(Q_{\mathrm{H}})^* \subseteq \bigcup_k \mathrm{FP}(I^*_{C^k}).$$

But $\mathrm{FP}(Q_{\mathrm{H}})^*$ is FP+C by Lemma 4.7. On the right-hand side of the above inclusion we apply Lemma 4.15 and the closure of FP with respect to interpretations to see that $\mathrm{FP}(I^*_{C^k}) \equiv \mathrm{FP}(I_{C^k})$. This proves $\mathrm{FP{+}C} \subseteq \bigcup_k \mathrm{FP}(I_{C^k})$. The converse inclusion follows directly from closure of FP+C with respect to interpretations (Proposition 4.8) and interpretability of I_{C^k} in FP+C (Lemma 4.14). □

The analogue of the Abiteboul-Vianu Theorem (Theorem 3.22) follows immediately.

Corollary 4.19. $\mathrm{FP{+}C} \equiv \mathrm{PFP{+}C}$ *if and only if* PTIME = PSPACE.

We may now also infer the basic inclusion $\mathrm{PFP{+}C} \subseteq C^{\omega}_{\infty\omega}$ from the characterization of PFP+C in Theorem 4.18 without getting involved in technicalities.

Corollary 4.20. $\mathrm{FP{+}C} \subseteq \mathrm{PFP{+}C} \subsetneq C^{\omega}_{\infty\omega}$

Proof. It suffices to show that every PFP-definable global relation is closed with respect to $\equiv^{C^k}$ for some k, cf. Lemma 1.33. But this is obvious from $\mathrm{PFP{+}C} \equiv \bigcup_k \mathrm{PFP}(I_{C^k})$. Strictness of the inclusion $\mathrm{PFP{+}C} \subsetneq C^{\omega}_{\infty\omega}$ is clear since PFP+C is in PSPACE whereas $C^{\omega}_{\infty\omega}$ expresses even non-recursive queries. □

There is of course also a straightforward direct proof of these inclusions parallel to the proof for FP, PFP $\subseteq L^{\omega}_{\infty\omega}$, cf. Lemma 1.29 and Corollary 1.30. Technically these are more tedious, however, since mixed-type predicates over the $\mathfrak{A}^*$ have to be represented in the one-sorted framework of the $\mathfrak{A}$ themselves. A single type $(1,1)$ formula $\varphi(x,\nu)$ of the two-sorted framework for instance can be decomposed into a family of formulae $\varphi_{n,j}(x)$ for $j \leqslant n$ with the intended meaning that for all $\mathfrak{A}$ of size n: $\varphi[\mathfrak{A}^*] = \bigcup_{0\leqslant j\leqslant n}\big(\varphi_{n,j}[\mathfrak{A}]\times\{j\}\big)$.

In characterizations like FP+C $\equiv$ PTIME$\big(I_{C^k}\big)$ for fixed points with counting, it is important to note that the size of I_{C^k} is of the same order as the size of the original structure. This essential difference between I_{L^k} and the I_{C^k} leads to a picture that is in sharp contrast with the second theorem of Abiteboul and Vianu for FP and PFP without counting (Theorem 3.24 above). Let PFP+C$|_{\text{poly}}$ be the sublogic of PFP+C in which all occurrences of the PFP-constructor must be such that the limit in the partial fixed-point process is always reached within a polynomial number of steps. The following very simple theorem shows FP+C to be better behaved as a logic for PTIME recursion within $C^{\omega}_{\infty\omega}$ than FP is within $L^{\omega}_{\infty\omega}$.

Theorem 4.21. PFP+C$|_{\text{poly}} \equiv$ FP+C.

Sketch of Proof. Let $\text{PFP}_{X,\overline{x\nu}}\varphi$ be such that the fixed-point process is polynomially bounded. This fixed-point process is then represented by a polynomially bounded PFP-process over the I_{C^k} for some k. Over the ordered I_{C^k} it must therefore be equivalent with an FP-process. Inductively we obtain PFP+C$|_{\text{poly}} \subseteq \bigcup_k \text{FP}(I_{C^k}) \equiv$ FP+C. □

4.3 The Separation from PTIME

It is an important result of Cai, Fürer and Immerman [CFI89] that also FP+C is too weak to capture the class of all PTIME queries on not necessarily ordered finite structures. The construction has been reviewed in Example 2.7.

It is worth to note that on the basis of the present analysis we may infer FP+C $\not\supseteq$ PTIME from the fact that none of the $C^k_{\infty\omega}$ defines *all* queries. This argument is of some interest in its own because it relativizes to many subclasses of the class of all finite structures. The only requirement on the subclass $\mathcal{K}$ is that it admits some kind of padding: some simple construction should be available within $\mathcal{K}$ that allows to increase arbitrarily the size of structures. We choose closure under disjoint unions as a corresponding prerequisite on $\mathcal{K}$ in the statement of the following theorem. It will be clear from the proof that a number of other natural closure conditions would serve just as well.

Theorem 4.22. *Let $\mathcal{K} \subseteq \text{fin}[\tau]$ be a class of finite τ-structures that is closed under disjoint unions. Assume that* FP+C *captures* PTIME *on $\mathcal{K}$, in particular that any* PTIME *computable boolean query on $\mathcal{K}$ is definable by a sentence* *of* FP+C *and hence also by a sentence in* $C^{\omega}_{\infty\omega}$. *Then there is some* k *satisfying the following two (equivalent) conditions.*

(i) I_{C^k} *classifies structures in* $\mathcal{K}$ *up to isomorphism: for all* $\mathfrak{A}, \mathfrak{A}' \in \mathcal{K}$: $I_{C^k}(\mathfrak{A}) = I_{C^k}(\mathfrak{A}') \iff \mathfrak{A} \simeq \mathfrak{A}'$.

(ii) In restriction to $\mathcal{K}$, $C^k_{\infty\omega} \equiv L_{\infty\omega}$ *for sentences; in other words, any boolean query on* $\mathcal{K}$ *must be definable in* $C^k_{\infty\omega}$.

Applying this to the class of all finite graphs, and using the result of Cai, Fürer and Immerman just to the effect that no $C^k_{\infty\omega}$ coincides with $L_{\infty\omega}$ on the class of all finite graphs (as expressed in Theorem 2.9) we obtain the following.

Corollary 4.23 (Cai, Fürer, Immerman).
FP+C $\subsetneq$ PTIME, *in fact even* PTIME $\not\subseteq C^{\omega}_{\infty\omega}$.

On the basis of Theorem 4.22 this separation is also obtained as a corollary of recent results of Gurevich and Shelah [GS96]. They prove that in a suitable vocabulary τ there are for each k rigid structures in fin[τ] that do not admit a $C^k_{\infty\omega}$-definable linear ordering. Again it follows that on (expansions) of these structures no $C^k_{\infty\omega}$ coincides with $L_{\infty\omega}$ (for sentences even).

Note that the only separation results between FP+C and PTIME that can be obtained along the lines of Theorem 4.22 — and these are all there are, as yet — are in fact separations of $C^{\omega}_{\infty\omega} \cap$ PTIME from PTIME.

Proof (of Theorem 4.22). Let $\mathcal{K}$ be as required. Choose some sufficiently fast growing monotone function $f: \omega \to \omega$ such that $f(n)$ is computable from n in time polynomial in $f(n)$. Assume that $f(n) > n$ for all n. It follows that there is a PTIME algorithm that recognizes numbers of the form $n\big(f(n) + f(n)^2\big)$ and computes n for these: for given m it suffices to compute $n\big(f(n) + f(n)^2\big)$ for all n with $n^3 \leqslant m$ and check for equality with m.

It further follows that m_1 and m_2 can be computed in PTIME from n and $m_1 f(n) + m_2 f(n)^2$ for any $m_1, m_2 \leqslant n$: simply expand the given number $m = m_1 f(n) + m_2 f(n)^2$ in base $f(n)$ to obtain the m_i as its digits.

We claim that for suitable f the following padded variant of the isomorphism query on $\mathcal{K}$ becomes a PTIME query:

$$Q := \Big\{ \mathfrak{C} \;\Big|\; \mathfrak{C} \simeq \underbrace{\mathfrak{A} \dot{\cup} \ldots \dot{\cup} \mathfrak{A}}_{m}, \text{ where } m = f(n) + f(n)^2, n = |A| \Big\}.$$

The intended algorithm first checks whether the size of an input $\mathfrak{C}$ is of the form $n\big(f(n) + f(n)^2\big)$ and computes n in this case. It then checks for all isomorphism types of connected[1] τ-structures $\mathfrak{D}$ of size at most n how many connected components of $\mathfrak{C}$ are isomorphic with $\mathfrak{D}$ (and that $\mathfrak{C}$ has no components of size greater than n). This is done in time polynomial in $|\mathfrak{C}|$

[1] A structure is called connected if it is not the disjoint union of two other structures.

provided $f(n)$ is sufficiently large for n; the precise meaning of 'sufficiently large' has to take into account the arities in τ.

Let $\nu(\mathfrak{D})$ be the corresponding number for each $\mathfrak{D}$. Then $\mathfrak{C} \in Q$ if and only if all $\nu(\mathfrak{D})$ are of the form $\nu(\mathfrak{D}) = \mu(\mathfrak{D})\big(f(n) + f(n)^2\big)$ for appropriate $\mu(\mathfrak{D}) \leqslant n$. Necessity of this condition is clear. For sufficiency observe that, if $\nu(\mathfrak{D}) = \mu(\mathfrak{D})\big(f(n) + f(n)^2\big)$ for all $\mathfrak{D}$, then $\mathfrak{C}$ is of the required form if for $\mathfrak{A}$ one takes the disjoint union of $\mu(\mathfrak{D})$ copies of each $\mathfrak{D}$.

By assumption Q therefore is definable in some $C^k_{\infty\omega}$. But the above characterization of $\mathfrak{C} \in Q$ through the $\nu(\mathfrak{D})$ also implies that for any two $\mathfrak{A}, \mathfrak{B} \in \mathrm{fin}[\tau]$ of the same size n,

$$\mathfrak{C} = \underbrace{\mathfrak{A} \,\dot\cup \ldots \dot\cup\, \mathfrak{A}}_{f(n)} \,\dot\cup\, \underbrace{\mathfrak{B} \,\dot\cup \ldots \dot\cup\, \mathfrak{B}}_{f(n)^2}$$

is in Q if and only if $\mathfrak{A} \simeq \mathfrak{B}$. It follows from Remark 4.16 on the other hand that $I_{C^k}(\mathfrak{C})$ is a function of $I_{C^k}(\mathfrak{A})$ and $I_{C^k}(\mathfrak{B})$, so that $\mathfrak{A} \simeq \mathfrak{B}$ is determined by $I_{C^k}(\mathfrak{A})$ and $I_{C^k}(\mathfrak{B})$. This implies claim (i) of the theorem, and equivalence of (i) and (ii) is obvious. The argument given here is a structural variant of the so-called *padding technique* that is often useful in complexity considerations. □

The results of this chapter show that FP+C is the right logic for PTIME recursion in the world of $C^\omega_{\infty\omega}$. In this respect its relation to $C^\omega_{\infty\omega}$ resembles that of FP to $L^\omega_{\infty\omega}$. It is known from the result of Cai, Fürer and Immerman that real PTIME is not within $C^\omega_{\infty\omega}$. On the other hand all known separation results for FP+C from PTIME are separations of $C^\omega_{\infty\omega} \cap$ PTIME from PTIME. The question that arises at this point is the following:

> Does FP+C capture PTIME $\cap\, C^\omega_{\infty\omega}$, the class of all those queries that are both PTIME computable and definable in $C^\omega_{\infty\omega}$?

More suggestively:

> Does FP+C capture PTIME in the world of $C^\omega_{\infty\omega}$?

This question is further explored in the last two chapters. Note that the same question with FP and $L^\omega_{\infty\omega}$ in the place of FP+C and $C^\omega_{\infty\omega}$ can be answered negatively unless PTIME = PSPACE. Obviously $\mathrm{PFP}|_{\mathrm{poly}} \subsetneq$ PTIME $\cap\, L^\omega_{\infty\omega}$, but $\mathrm{PFP}|_{\mathrm{poly}} \subseteq$ FP only if PTIME = PSPACE by the second theorem of Abiteboul and Vianu. There is a reasonable variant of the issue that remains an open problem for FP and $L^\omega_{\infty\omega}$, too. We shall come back to these issues in Chapter 6. In the last Chapter we find positive solutions to such questions in the very restricted case of just two variables, i.e. for $L^2_{\infty\omega}$ and $C^2_{\infty\omega}$.

4.4 Other Characterizations of FP+C

It may be a further indication of the naturalness of FP+C as a level of expressiveness within PTIME that it admits several different equivalent logical characterizations and also a natural algorithmic characterization. We here only indicate some of these briefly. More detailed accounts can be found in [GO93] and [Ott96a], respectively.

Among the logical variations we mention the following:

(a) FP+C can be obtained as a straightforward extension of Datalog. For our purposes Datalog is the logic of positive Horn-clause programs with the least fixed-point semantics. Its counting extension is based on the two-sorted variants of structures in $\mathrm{fin}[\tau]^*$ and allows the use of counting terms and cardinality comparisons in the sense of $\leqslant$ in clauses. It is not difficult to see that the counting extension leads to closure under negation. It follows that this extension of Datalog comprises the full power of fixed-point logic and thus is semantically equivalent with FP+C.

(b) The approach to extend finite structures with standard sorts, like the arithmetical second sort of the structures in $\mathrm{fin}[\tau]^*$, has been carried much further in the framework of *meta-finite structures* put forward by Grädel and Gurevich in [GG95]. Here finiteness of the second standardized sorts is given up in order to obtain a more uniform modelling for issues on finite structures that essentially involve reference to infinite standard structures (like the natural or the real numbers). In order to obtain an adequate limitation on the access to the infinite standard domains, recursive processes like those in fixed-point are restricted to the finite relational domain. The infinite standard parts are accessed through terms and *multiset operations.* The latter can roughly be described as arithmetical operations that are performed on weight functions from the finite relational domain to the infinite standard part. It turns out that FP+C can be isolated in this framework by taking arithmetic on the natural numbers $(\omega, <, +, \cdot)$ for the infinite standard structure, with exactly the PTIME multiset operations. It is shown in [GG95] that the expressive power of fixed-point logic in this meta-finite frame coincides with FP+C.

We mention two more characterizations of different kinds in slightly greater detail. One is in terms of uniform sequences of formulae, the other by means of a computational model.

P-uniform sequences of formulae. Logical characterizations in terms of sequences of formulae are proposed and investigated in the work of Immerman, see for instance [Imm82]. Let $\mathrm{fin}_n[\tau]$ stand for the restriction of $\mathrm{fin}[\tau]$ to structures of size n. The idea is to associate for instance with a boolean query $Q \subseteq \mathrm{fin}[\tau]$ a *sequence* of sentences $\left(\varphi_n\right)_{n\geqslant 1}$ in some logic $\mathcal{L}[\tau]$ such that for all sizes n:

$$Q \cap \mathrm{fin}_n[\tau] = \left\{\mathfrak{A} \in \mathrm{fin}_n[\tau] \,\middle|\, \mathfrak{A} \models \varphi_n\right\}.$$

A priori this is a completely non-uniform notion of logical definability. Restrictions on the constituent formulae φ_n in terms of quantifier rank, numbers of variables and size (all regarded as functions in n) or constructibility criteria for the mapping $n \mapsto \varphi_n$ serve to employ this approach as a tool in the logical analysis of complexity. It turns out that FP+C and PFP+C are isolated by very natural uniformity conditions on sequences. Note that in the presence of counting quantifiers and for sequences of formulae $\varphi_n \in C^k_{\omega\omega}$ the semantics given to the sequence is that of $\bigvee_{n \geqslant 1} (\exists^{=n} x\, x = x \wedge \varphi_n) \in C^k_{\infty\omega}$.

Definition 4.24. *Call a sequence $(\varphi_n)_{n \geqslant 1}$ of formulae in some $C^k_{\omega\omega}$* PTIME-uniform, *respectively* PSPACE-uniform, *if φ_n is constructible in time, respectively space, polynomial in n. Let* PTIME–$C^\omega_{\infty\omega}$ *and* PSPACE–$C^\omega_{\infty\omega}$ *stand for the sublogics of $C^\omega_{\infty\omega}$ corresponding to all* PTIME-, *respectively* PSPACE-*uniform sequences.*

Clearly PTIME–$C^\omega_{\infty\omega} \subseteq$ PSPACE–$C^\omega_{\infty\omega} \subseteq C^\omega_{\infty\omega}$. The following is proved in [Ott96a].

Proposition 4.25. FP+C $\equiv$ PTIME–$C^\omega_{\infty\omega}$ *and* PFP+C $\equiv$ PSPACE–$C^\omega_{\infty\omega}$.

This is quite unlike the situation for FP and PFP themselves: trivial examples involving pure sets show for instance that FP is properly contained in the correspondingly defined PTIME–$L^\omega_{\infty\omega}$.

A computational characterization. Finally there is a natural computational model whose PTIME and PSPACE restrictions coincide with FP+C and PFP+C, respectively. This model is the obvious generalization of the relational computational model of Abiteboul and Vianu [AV91] that incorporates counting operations in a generic manner. Let us call the machines under consideration *relational machines with counting.* We give a brief sketch. A relational machine with counting consists of two components. First there is a *relational store* with a fixed number of relational registers of fixed arities. These can hold sets of tuples from the domain of the input structure. Among these relational registers there are specified ones that are initialized to represent the given predicates in the input structure. The others are initially empty. In any case, at each stage of the computation, the content of a relational register is a relation over the input domain. The second component of the machine resembles an ordinary Turing machine with a work tape with a read–write head, an extra communication tape with a write–only head, and the usual finite state control. The interaction between the two components is the following.

- Each transition, as laid down in the transition table of the Turing control, may depend not only on the current internal state and symbol read on the work tape but also on the information which of the relational registers are currently empty. These implicit emptiness queries constitute the only flow of information from the relational part of the machine to the Turing component.
- The execution of a transition may involve not only the printing of tape symbols and movements of the heads but may also include one of several

update operations on the relational store. The operations available here are the following:

- copy and move operations between relational registers.
- boolean operations on the current contents of specified relational registers (e.g. union and complementation).
- operations corresponding to the natural action of the permutation groups S_r on the contents of r-ary registers.
- counting projections.

Counting projections take as input a numerical parameter ν whose current value is read from the communication tape. The content of a prescribed relational register R is then replaced by all those tuples for which there are at least ν substitutes for the first component that are currently in R: $R' := \{\overline{a} \mid \exists^{\geqslant\nu} b(\overline{a}\frac{b}{1} \in R)\}$. Only in this operation does the present model extend the one proposed by Abiteboul and Vianu. Their model only allows ordinary existential projections which appear here as a special case for $\nu = 1$.

Computations of these machines are formalized in the natural manner. The result of a computation that is to produce a boolean value can be encoded in the final state reached by the Turing control. For machines that are to compute an r-ary query, the output is the content of one specified r-ary relational register when the machine reaches its halting state.

This model of computation is entirely isomorphism-preserving ('generic' is the term usually applied in the literature). Any isomorphism between input structures naturally extends to all stages of the computation, so that the resulting computations are not only equivalent but really isomorphic themselves.

Complexities for this model are defined in terms of the Turing component. PTIME and PSPACE for the relational machines with counting comprise those queries that are computable by one of these machines within a number of steps, respectively with the use of a number of tape cells (of the Turing component) that is polynomially bounded in the size of the input structure.

Theorem 4.26. *On finite relational structures,* PTIME *and* PSPACE *for the relational machines with counting exactly correspond to* FP+C *and* PFP+C.

In particular FP+C and PFP+C are the polynomial time and space restrictions of a generic model of computation — a situation that in a sense is ruled out by the second theorem of Abiteboul and Vianu, Theorem 3.24 above, for FP and PFP themselves.

Attributions and remarks. FP+C — roughly in our formalization — is implicit in the work of Immerman, in particular see [Imm87a]. The present explicit form was first presented in [GO93]. Most of the material treated in this chapter can also be found in [Ott96a] and [GO93]. The latter source should be consulted in particular for those characterizations of FP+C that are only sketched here.