3. The Invariants

We introduce *complete structural invariants* that classify finite relational structures up to C^k- and L^k-equivalence, respectively. These invariants are based on the definable pre-orderings with respect to C^k- and L^k-types obtained in the analysis of the games in the preceding chapter. The invariants are PTIME computable and inherit specific definability properties from the pre-orderings with respect to types. These definability properties and a close relationship with the fixed-point logics make the invariants extremely useful in investigations concerning fixed-point logics and complexity issues. This approach has been initiated and lead to success in the seminal work of Abiteboul and Vianu. They first introduced a kind of ordered invariants with respect to their model of relational computation and with this technique derived important results concerning the relationship between FP and PFP.

• In the introductory Section 3.1 we relate the concept of the proposed invariants to the abstract notion of complete invariants.

• Section 3.2 provides the definition of our C^k-invariants and states their fundamental definability properties.

• Section 3.3 similarly treats the invariants for L^k.

• In Section 3.4 we consider applications of the invariants to the analysis of fixed-point logics. A main point is the discussion of the Abiteboul-Vianu Theorem on the relation between FP and PFP. As far as the C^k-invariants are concerned, the corresponding considerations are of a preliminary nature here. This analysis will be pursued further in Chapter 4 where it becomes possible to link the C^k-invariants directly with fixed-point logic with counting. We include here a comparison between the C^k- and the L^k-invariants.

• In Section 3.5 it is indicated that — up to interpretability in powers — our invariants essentially reduce to the two-dimensional ones, i.e. to those for C^2 and L^2.

3.1 Complete Invariants for L^k and C^k

Recall from Definition 1.58 the notion of a complete invariant for an equivalence relation $\sim$: I is a complete invariant for $\sim$ if I classifies objects exactly up to $\sim$: $x \sim x'$ if and only if $I(x) = I(x')$. We apply this notion to the equivalence relations $\equiv^{\mathcal{L}}$, $\mathcal{L} = C^k_{\infty\omega}$ or $L^k_{\infty\omega}$. These may be regarded as equivalence relations on fin$[\tau]$ as well as on the fin$[\tau; r]$ for $r \leqslant k$. Accordingly we actually get two notions of complete invariants in each case.

Definition 3.1. *Let $\mathcal{L}$ be a logic. A functor I on* fin$[\tau]$ *is a* complete $\mathcal{L}$-invariant *if*

$$\forall \mathfrak{A} \forall \mathfrak{A}' \quad \mathfrak{A} \equiv^{\mathcal{L}} \mathfrak{A}' \iff I(\mathfrak{A}) = I(\mathfrak{A}').$$

Similarly, I is a complete invariant for $\mathcal{L}$ on fin$[\tau; r]$ *if for $(\mathfrak{A}, \overline{a})$ and $(\mathfrak{A}', \overline{a}')$ in* fin$[\tau; r]$*: $(\mathfrak{A}, \overline{a}) \equiv^{\mathcal{L}} (\mathfrak{A}', \overline{a}') \Leftrightarrow I(\mathfrak{A}, \overline{a}) = I(\mathfrak{A}', \overline{a}')$.*

A computable complete invariant I is a corresponding mapping from fin$[\tau]$ or fin$[\tau; k]$ to some set S of objects with a standard encoding realized by an algorithm $\mathcal{A}$: stan$[\tau] \to S$ or $\mathcal{A}$: stan$[\tau; k] \to S$. Compare Definition 1.61. The C^k- and L^k-invariants considered in the following are PTIME computable and take as their values linearly ordered structures (or structures over standard domains n). One of the goals of this chapter is the following theorem.

Theorem 3.2. *There are* PTIME *computable complete $\mathcal{L}$-invariants for $\mathcal{L} = C^k_{\infty\omega}$ and $L^k_{\infty\omega}$.*

The backbones of the invariants are the ordered representations of the $A^k / \equiv^{\mathcal{L}}$ derived in the preceding chapter. On fin$[\tau]$ we shall have

$$I_{\mathcal{L}}(\mathfrak{A}) = \left(A^k \big/ \equiv^{\mathcal{L}}, \leqslant, \underbrace{\ldots\ldots}_{(*)} \right),$$

where $\leqslant$ is the linear ordering of the quotient that is interpreted by the corresponding pre-ordering with respect to types over A^k. $(*)$ encodes additional combinatorial information so that exactly the structural information for $\mathcal{L}$-games over $\mathfrak{A}$ is retrievable from $I_{\mathcal{L}}(\mathfrak{A})$. This ensures that $I_{\mathcal{L}}(\mathfrak{A})$ comprises a complete description of the $\mathcal{L}$-theory of $\mathfrak{A}$ as required of a complete $\mathcal{L}$-invariant over fin$[\tau]$.

By definition the rôle of the invariants over fin$[\tau]$ is comparable to that of Scott sentences — they provide concise abstractions of the complete theories of structures. While Scott sentences may be regarded as the syntactic correlate of the games, the proposed invariants are structural correlates of the games. This structural aspect of the invariants has particular advantages in the further model theoretic applications: these invariants are adapted to simulate fixed-point processes over the original structures in a natural manner as we shall see in Section 3.4.

3.2 The C^k-Invariants

We introduce and discuss the C^k-invariants on fin[τ]. Complete invariants for C^k-equivalence on fin[$\tau; r$] for $r \leqslant k$ are easily derived as extensions of those on fin[τ]. The invariants are built upon the ordered quotients $(A^k/\equiv^{C^k}, \leqslant)$, where $\leqslant$ is the ordering induced by the pre-ordering $\preccurlyeq$ with respect to C^k-types, compare Theorem 2.28. In order to put the full information about the C^k-game over $\mathfrak{A}$ — or about the complete C^k-theory of $\mathfrak{A}$ — into the invariant, this ordered quotient is expanded with the following:

(i) Atomic components of types:
for each atomic type $\theta \in \mathrm{Atp}(\tau; k)$ the unary predicate P_θ is introduced on $A^k/\equiv^{C^k}$. For $\alpha \in A^k/\equiv^{C^k}$ put $\alpha \in P_\theta$ if $\mathrm{atp}_{\mathfrak{A}}(\overline{a}) = \theta$ for $\overline{a} \in \alpha$.

(ii) Accessibility:
for each j the binary predicate E_j, which encodes accessibility in moves concerning the j-th component, is transferred to $A^k/\equiv^{C^k}$ as follows. For $\alpha_1, \alpha_2 \in A^k/\equiv^{C^k}$ put $(\alpha_1, \alpha_2) \in E_j$ if for $\overline{a} \in \alpha_1$ there is some $b \in A$ with $\overline{a}\frac{b}{j} \in \alpha_2$.

(iii) Symmetries:
for each permutation ρ in the symmetric group S_k acting on $\{1, \ldots, k\}$ a binary predicate S_ρ is defined on $A^k/\equiv^{C^k}$ by:
$(\alpha_1, \alpha_2) \in S_\rho$ if $\rho(\overline{a}) \in \alpha_2$ for $\overline{a} \in \alpha_1$.

(iv) Multiplicities:
for each j a weight function ν_j from $A^k/\equiv^{C^k}$ to natural numbers is introduced which sends α to $\left|\left\{b \in A \mid \overline{a}\frac{b}{j} \in \alpha\right\}\right|$, for $\overline{a} \in \alpha$.

It has to be checked that the given definitions are independent of choices of representatives for the $\equiv^{C^k}$-classes on A^k. Recall that $A^k/\equiv^{C^k} = \mathrm{Tp}^{C^k}(\mathfrak{A}; k)$. Clearly for the P_θ, $\alpha \in P_\theta$ if $\alpha \models \theta$. For the others choose for each C^k-type α a $C^k_{\infty\omega}$-formula $\varphi_\alpha(\overline{x})$ that isolates α (cf. Lemma 1.33). Consider now the E_j. For any two C^k-types $\alpha, \beta \in \mathrm{Tp}^{C^k}(\mathfrak{A}; k)$, either $\alpha \models \exists x_j \varphi_\beta(\overline{x})$ or $\alpha \models \neg\exists x_j \varphi_\beta(\overline{x})$. Accessibility of a position of type β in the j-th component is thus determined by the type α of the given position. Similarly for the multiplicities ν_j. For each α there must be some natural number m such that $\alpha \models \exists^{=m} x_j \varphi_\alpha(\overline{x})$; this m is the value of $\nu_j(\alpha)$. The operation of $\rho \in S_k$ preserves C^k-equivalence so that the representation of ρ as a binary predicate on the quotient is also sound. Alternatively the operations ρ may actually be defined syntactically on C^k-types through an operation on the variable symbols.

Note that the information about atomic types, accessibility and permutations are encoded by relations over $A^k/\equiv^{C^k}$ whereas for the multiplicities we have to resort to external weight functions. The values of the ν_j over $A^k/\equiv^{C^k}$ are bounded by $|A|$.

Definition 3.3. *Let for each k and each fixed finite relational vocabulary τ, the C^k-invariant I_{C^k} be the functor which sends a finite τ-structure $\mathfrak{A}$ to the weighted linearly ordered structure*

$$I_{C^k}(\mathfrak{A}) = \Big(A^k/{\equiv^{C^k}}, \leqslant, (P_\theta), (E_j), (S_\rho); (\nu_j)\Big),$$

where θ ranges over $\mathrm{Atp}(\tau; k)$, *j over $\{1, \dots, k\}$, and ρ over S_k. The interpretations of $\leqslant$, the P_θ, E_j, S_ρ and ν_j are as defined above.*

Obviously I_{C^k} is a PTIME computable functor. Formally we regard the *relational part* of $I_{C^k}(\mathfrak{A})$,

$$(A^k/{\equiv^{C^k}}, \leqslant, (P_\theta), (E_j), (S_\rho)),$$

as a relational structure on the standard universe of size $|A^k/\equiv^{C^k}| \leqslant |A|^k$ with its natural ordering. The weight functions ν_j take values in $\{1, \dots, |A|\}$. $I_{C^k}(\mathfrak{A})$ as a whole may therefore in some canonical way be encoded as a relational structure over the standard universe of size $|A|$, if k-tuples are used to encode the elements of $A^k/\equiv^{C^k}$.

Remark 3.4. *We regard I_{C^k} as a* PTIME *functor that takes standard objects — namely canonical relational encodings of the weighted ordered structures $I_{C^k}(\mathfrak{A})$ over standard domains of size $|A|$ — as its values. The* size *of I_{C^k} is taken to be $|A|$.*[1]

The data encoded in $I_{C^k}(\mathfrak{A})$ are redundant in several respects. In the presence of the $\rho \in S_k$ it would in particular suffice to keep only one of the E_j and only one of the ν_j. For instance $\nu_{j_1} = \nu_{j_2} \circ \rho$ where ρ is the permutation exchanging j_1 and j_2. We keep this redundancy because the highly symmetric format is easier to handle in some applications.

It is slightly less obvious that also the S_ρ are PTIME computable (and hence FP-definable) from the complete set of the E_j and ν_j. To see this observe that the quotients $A^k/\approx_i$, that occur in the inductive generation of $\approx$, are all naturally interpreted over $A^k/\equiv^{C^k}$: the equivalence classes of the $\approx_i$ are unions of $\equiv^{C^k}$-classes, as $\equiv^{C^k}$ is a refinement of the $\approx$. At $\approx_0$-level, the classes are just the P_θ over $A^k/\equiv^{C^k}$. The operation of S_k on $\mathrm{Atp}(\mathfrak{A}; k) = A^k/\approx_0$ is trivially definable in this interpretation. $\rho \in S_k$ sends P_θ to $P_{\rho(\theta)}$, where $\rho(\theta)$ is obtained by operating with ρ^{-1} on the variables in θ. Inductively, in each refinement step, the operation of S_k on $A^k/\approx_{i+1}$ is determined by that on $A^k/\approx_i$. The refinement is governed by the values of the functions $\nu_j^\alpha(\overline{a}) = \big|\{b \in A \;/\; \overline{a}\frac{b}{j} \in \alpha\}\big|$ for $\alpha \in A^k/\approx_i$. And for these we obviously have

[1] The value $|A|$, rather than for instance $|A|^k$ or $\sum_j \sum_\alpha \nu_j(\alpha)$ is a matter of convention. The size of $I_{C^k}(\mathfrak{A})$ is naturally only determined up to polynomial transformations. The point is that the size of $I_{C^k}(\mathfrak{A})$, with weights taken into account, is polynomially related to the size of $\mathfrak{A}$, and not to $|A^k/{\equiv^{C^k}}|$.

$$\nu_j^\alpha(\overline{a}) = \nu_{\rho(j)}^{\rho(\alpha)}\big(\rho(\overline{a})\big) \quad \text{for all } \rho \in S_k.$$

Furthermore, even the ordering $\leqslant$ is PTIME computable form the remaining data on $A^k/\equiv^{C^k}$, since the entire refinement process in the generation of the $\preccurlyeq_i$ can also be simulated over $\big(A^k/\equiv^{C^k}, (P_\theta), (E_j); (\nu_j)\big)$.

These facts are stated for future reference in the following remark.

Remark 3.5. *The ordering $\leqslant$ and the interpretations of the S_ρ in $I_{C^k}(\mathfrak{A})$ are* PTIME *computable from the reduct of the I_{C^k} to vocabulary consisting only of the E_j and the P_θ together with the weight functions ν_j.*

The relational part of the I_{C^k} gets naturally interpreted over the original structures as a quotient over the k-th power. By definition it *is* the quotient of the k-th power of the universe with respect to $\equiv^{C^k}$. More precisely, we get the following. Recall that $\mathrm{FP}(Q_\mathrm{R})$ is fixed-point logic with the Rescher quantifier.

Proposition 3.6. *The relational part of $I_{C^k}(\mathfrak{A})$, i.e. the relational structure $\big(A^k/\equiv^{C^k}, \leqslant, (P_\theta), (E_j), (S_\rho)\big)$ is $\mathrm{FP}(Q_\mathrm{R})$-interpretable as a quotient over the k-th power of $\mathfrak{A}$. Moreover, the weights ν_j are the cardinalities of atomically definable predicates in this interpretation.*

Sketch of Proof. The intended interpretation is straightforward since I_{C^k} is defined as a quotient on the k-th power. $\mathrm{FP}(Q_\mathrm{R})$-definability of the relational part is also obvious. $\mathrm{FP}(Q_\mathrm{R})$ is needed to define the equivalence relation $\equiv^{C^k}$ for the quotient and for the linear ordering $\leqslant$ on this quotient; this is just $\mathrm{FP}(Q_\mathrm{R})$-definability of the pre-ordering $\preccurlyeq$ as stated in Theorem 2.28. The P_θ, E_j and S_ρ are in fact first-order interpretable relative to the interpreted $\big(A^k/\equiv^{C^k}, \leqslant\big)$. The ν_j finally are defined in terms of this interpretation over A^k according to $\nu_j(\alpha) = \big|\{\overline{b} \in \alpha | E_j\overline{a}\overline{b}\}\big|$ for any $\overline{a} \in \alpha$. □

It remains to establish the I_{C^k} as complete invariants for C^k on $\mathrm{fin}[\tau]$.

Theorem 3.7. *The functor I_{C^k} is a complete $C^k_{\infty\omega}$-invariant on* $\mathrm{fin}[\tau]$. *It classifies finite τ-structures exactly up to equivalence in C^k:*

$$\forall\mathfrak{A}\forall\mathfrak{A}' \quad I_{C^k}(\mathfrak{A}) = I_{C^k}(\mathfrak{A}') \iff \mathfrak{A} \equiv^{C^k} \mathfrak{A}'.$$

Proof. By Lemma 1.34 $\mathfrak{A} \equiv^{C^k} \mathfrak{A}'$ if and only if $\mathfrak{A}$ and $\mathfrak{A}'$ realize exactly the same C^k-types. This is used in the proof.

i) Assume first that $\mathfrak{A} \equiv^{C^k} \mathfrak{A}'$, so that $\mathfrak{A}$ and $\mathfrak{A}'$ realize exactly the same C^k-types. The crucial fact for the proof that $I_{C^k}(\mathfrak{A}) = I_{C^k}(\mathfrak{A}')$ is that these types get ordered in exactly the same way by $\preccurlyeq$ over $\mathfrak{A}$ and $\mathfrak{A}'$. This follows from the global view of $\preccurlyeq$ on $\mathrm{fin}[\tau; k]$ as expressed in Lemma 2.29 and Corollary 2.30. It follows that the natural isomorphism between $\big(A^k/\approx, \leqslant\big)$ and $\big(A'^k/\approx, \leqslant\big)$ as ordered structures is the identity function on C^k-types. Thus $I_{C^k}(\mathfrak{A}) = I_{C^k}(\mathfrak{A}')$ follows directly from the definitions, since we have seen

above that all the extra information encoded in the invariants is determined by the constituent types.
ii) For the converse implication it suffices to prove that the C^k-type corresponding to an element $\alpha \in I_{C^k}(\mathfrak{A})$ can be recovered from the invariant. Recall that the universe $A^k/\equiv^{C^k}$ of $I_{C^k}(\mathfrak{A})$ is the set $\mathrm{Tp}^{C^k}(\mathfrak{A};k)$ of C^k-types realized over $\mathfrak{A}$. The claim is clear at the atomic level because of the unary predicates P_θ. Inductively assume that for each formula $\varphi(\overline{x}) \in C^k_{\infty\omega}$ of quantifier rank at most i the subset $\underline{\varphi} := \{\alpha \in A^k/\equiv^{C^k} \mid \varphi \in \alpha\}$ has been determined as a subset of $I_{C^k}(\mathfrak{A})$. Without loss of generality consider a formula $\exists^{\geqslant m} x_j \varphi(\overline{x})$ with φ of quantifier rank at most i for the inductive step. It follows from the definition of the E_j and ν_j that $(\exists^{\geqslant m} x_j \varphi) \in \alpha$ if and only if

$$\sum_{(\alpha,\alpha')\in E_j,\alpha'\in\underline{\varphi}} \nu_j(\alpha') \geqslant m .$$

Therefore $I_{C^k}(\mathfrak{A})$ fully determines the set of C^k-types realized in $\mathfrak{A}$, and thus the complete C^k-theory of $\mathfrak{A}$. □

The proof also shows that classification up to $\equiv^{C^k}$ naturally extends from structures in $\mathrm{fin}[\tau]$ to structures with parameters, in particular to the classification of $\mathrm{fin}[\tau;k]$ up to $\equiv^{C^k}$. This is expressed in the following corollary.

Corollary 3.8. *The following extension of the I_{C^k} to* $\mathrm{fin}[\tau;k]$ *provides a complete invariant for C^k on* $\mathrm{fin}[\tau;k]$*:*

$$(\mathfrak{A},\overline{a}) \longmapsto \Big(I_{C^k}(\mathfrak{A}),[\overline{a}]\Big),$$

where $[\overline{a}]$ is that element of $I_{C^k}(\mathfrak{A})$ representing $\mathrm{tp}^{C^k}_{\mathfrak{A}}(\overline{a})$. In terms of the interpretation of the relational part of $I_{C^k}(\mathfrak{A})$ over $\mathfrak{A}$ it is just the equivalence class of $\overline{a}$ as induced by $\preccurlyeq$. The extended invariants share all the above-mentioned definability and interpretability properties with the standard ones.

Embedding A^r into A^k and $\mathrm{Tp}^{\mathcal{L}}(\mathfrak{A};r)$ into $\mathrm{Tp}^{\mathcal{L}}(\mathfrak{A};k)$ for $r \leqslant k$ via $(a_1,\ldots,a_r) \mapsto (a_1,\ldots,a_1,a_1,\ldots,a_r)$, with $k-r$ additional entries of a_1 as usual, we similarly obtain complete invariants for C^k on $\mathrm{fin}[\tau;r]$ *for all $r \in \{1,\ldots,k\}$.*

Corollary 3.9. *C^k-equivalence on* $\mathrm{fin}[\tau]$ *as well as on* $\mathrm{fin}[\tau;k]$ *is in* PTIME. *For any finite relational τ there are* PTIME *algorithms that decide for $\mathfrak{A}$ and $\mathfrak{A}'$ in* $\mathrm{fin}[\tau]$ *whether $\mathfrak{A} \equiv^{C^k} \mathfrak{A}'$ and for $(\mathfrak{A},\overline{a})$ and $(\mathfrak{A}',\overline{a}')$ in* $\mathrm{fin}[\tau;k]$ *whether $(\mathfrak{A},\overline{a}) \equiv^{C^k} (\mathfrak{A}',\overline{a}')$.*

These algorithms need merely evaluate the corresponding invariants and check for equality. C^k-equivalence of finite relational structures can also be shown to be definable in the extension of fixed-point logic by the Härtig quantifier — when suitably formalized as a query C^k-EQ.

Let to this end τ' consist of τ together with two new unary predicate symbols U_1 and U_2. Pairs of structures in fin[τ] are naturally encoded as τ'-structures in which the U_i separate the universe into two disjoint subsets for the universes of the individual τ-structures. Correspondingly put

$$C^k\text{-EQ} = \Big\{(\mathfrak{B}, U_1, U_2) \in \text{fin}[\tau \dot{\cup} \{U_1, U_2\}] \Big| B = U_1 \dot{\cup} U_2, \mathfrak{B} \upharpoonright U_1 \equiv^{C^k} \mathfrak{B} \upharpoonright U_2\Big\}.$$

It is obvious that this class is definable in FP(Q_{R}), fixed-point logic with the Rescher quantifier. FP(Q_{R}) interprets the relational parts of the I_{C^k} of the individual structures over the pair structure. Whether two interpreted linearly ordered relational structures are isomorphic is even FP-definable. For agreement also in the weight functions it suffices to check equalities of the cardinalities that define their values, and this can be done in a fixed-point process that uses the Rescher quantifier for the individual cardinality comparisons.

Now this can be strengthened to definability in the weaker extension of FP by the Härtig quantifier, for cardinality equality (compare Definition 1.53). In fact the invariants themselves as ordered structures need not actually be evaluated. We may instead directly consider the inductive generation of the stages $\approx_i$ in the generation of $\approx \, = \, \equiv^{C^k}$ over positions in $\mathfrak{A}$ and $\mathfrak{A}'$.

Checks for cardinality equality suffice for a fixed-point process whose stages are the $\approx_i$ in restriction to $\mathfrak{A}$ and $\mathfrak{A}'$. Compare Proposition 2.15 where the refinement step in the inductive definition of the $\approx_i$ is formalized in terms of cardinality equalities. (Formally one should use the complements of the $\approx_i$ to make the inductive process increasing.)

$\mathfrak{A} \equiv^{C^k} \mathfrak{A}'$ if the restriction of $\approx$ to $A^k \times A'^k$ induces an isomorphism between $\big(A^k/\approx, (P_\theta), (E_j); (\nu_j)\big)$ and $\big(A'^k/\approx, (P_\theta), (E_j); (\nu_j)\big)$. As these quotients and the candidate isomorphism between them are FP(Q_{H})-interpretable, FP(Q_{H}) also suffices to check the isomorphism property: Q_{H} is here used again for equality checks on the weights. We thus get the following.

Proposition 3.10. *C^k-equivalence of structures is definable in* FP(Q_{H}), *fixed-point logic with the Härtig quantifier, in the sense that the above class C^k-EQ of encodings of pairs of C^k-equivalent structures is definable in* FP(Q_{H}).

3.3 The L^k-Invariants

We sketch the introduction of the corresponding complete invariants I_{L^k} for the L^k. These are based on the pre-orderings with respect to L^k-type as characterized in Theorem 2.31 of the preceding chapter. The ordered quotients $\big(A^k/\equiv^{L^k}, \leqslant\big)$ are augmented by exactly the same relational expansions as in the case of I_{C^k}. There are, of course, no numerical weights to be kept here.

Definition 3.11. *For each k and each fixed finite relational vocabulary τ, the L^k-invariant I_{L^k} is the following* PTIME *computable functor which sends a finite τ-structure $\mathfrak{A}$ to the linearly ordered structure*

$$I_{L^k}(\mathfrak{A}) = \left(A^k/{\equiv^{L^k}}, \leqslant, (P_\theta), (E_j), (S_\rho)\right),$$

where θ ranges over $\mathrm{Atp}(\tau; k)$, *j over $\{1, \ldots, k\}$, and ρ over S_k.*

$\leqslant$ is the linear ordering induced by the pre-ordering with respect to L^k-types (Theorem 2.31), the P_θ, E_j and S_ρ are defined exactly as for I_{C^k}:

- *P_θ contains those $\alpha \in A^k/\equiv^{L^k} = \mathrm{Tp}^{L^k}(\mathfrak{A}; k)$, for which $\alpha \models \theta$;*
- *$(\alpha_1, \alpha_2) \in E_j$ if for $\overline{a} \in \alpha_1$ there is some $b \in A$ such that $\overline{a}\frac{b}{j} \in \alpha_2$;*
- *$(\alpha_1, \alpha_2) \in S_\rho$ if for $\overline{a} \in \alpha_1$ the permuted tuple $\rho(\overline{a})$ is in α_2.*

Formally $I_{L^k}(\mathfrak{A})$ is regarded as a relational structure on the standard universe of size $\left|A^k/\equiv^{L^k}\right|$ with its natural ordering.

As with the I_{C^k} above it would suffice to keep one of the E_j because the others are definable from any particular one with the help of the S_ρ. And again, in the presence of all the E_j, the S_ρ and the ordering $\leqslant$ are PTIME computable from the remaining data in I_{L^k} (compare Remark 3.5).

In analogy with Proposition 3.6 for the I_{C^k} we here obtain the following.

Proposition 3.12. *$I_{L^k}(\mathfrak{A})$ is* FP*-interpretable as a quotient over the k-th power of $\mathfrak{A}$.*

And of course the I_{L^k} are complete invariants for L^k.

Theorem 3.13. *The functor I_{L^k} is a complete invariant for $L^k_{\infty\omega}$ on* $\mathrm{fin}[\tau]$. *It classifies finite τ-structures exactly up to equivalence in L^k:*

$$\forall\mathfrak{A}\forall\mathfrak{A}' \quad I_{L^k}(\mathfrak{A}) = I_{L^k}(\mathfrak{A}') \iff \mathfrak{A} \equiv^{L^k} \mathfrak{A}'.$$

The proof can be given along exactly the same lines as that for Theorem 3.7 with the obvious simplifications. Extensions of the I_{L^k} to complete invariants on the $\mathrm{fin}[\tau; r]$ for $r \leqslant k$ are obtained as in Corollary 3.8.

Corollary 3.14. *L^k-equivalence over* $\mathrm{fin}[\tau]$ *as well as over* $\mathrm{fin}[\tau; k]$ *can be checked through evaluation of the corresponding invariants and hence is in* PTIME.

We get more, namely FP-definability of L^k-equivalence as a query on pair structures, a result due to Kolaitis and Vardi [KVa92b]. Putting

$$L^k\text{-EQ} = \left\{(\mathfrak{B}, U_1, U_2) \in \mathrm{fin}[\tau\dot{\cup}\{U_1, U_2\}] \;\middle|\; B = U_1 \dot{\cup} U_2, \mathfrak{B} \restriction U_1 \equiv^{L^k} \mathfrak{B} \restriction U_2\right\},$$

we obtain the following corollary.

Corollary 3.15 (Kolaitis, Vardi). *L^k-equivalence is* FP*-definable in the sense that the class L^k-EQ is definable in* FP. *In particular L^k-equivalence is in* PTIME.

Sketch of Proof. The invariants I_{L^k} are FP-interpretable and isomorphism of embedded linearly ordered structures is obviously in FP. The claim follows with closure of FP under generalized interpretations. □

3.4 Some Applications

3.4.1 Fixed-Points and the Invariants

From Lemma 1.29 we know that the fixed-point logics FP and PFP are sublogics of $L^{\omega}_{\infty\omega}$. For any fixed-point process there is some k such that this fixed-point process and its stages are all $L^k_{\infty\omega}$-definable. But this implies in particular that this fixed-point process does not distinguish between $L^k_{\infty\omega}$-equivalent tuples. In other words, the generation of the fixed point on $\mathfrak{A}$ can faithfully be represented on the quotient $A^k/\equiv^{L^k}$. This observation is the key to important insights into the nature of FP and PFP in relation to computational complexity that are due to Abiteboul and Vianu. In this first part we present the technical basis.

Let $\varphi(Z_1, \ldots, Z_l, \overline{x}) \in L^k_{\omega\omega}[\tau]$ be free in the indicated variables. Assume that the arity of the Z_i is at most k and that $\overline{x}$ is a tuple of (at most k) distinct variables. We want show that there is a first-order formula $\underline{\varphi}$ that captures the semantics of φ over the quotients $A^k/\equiv^{L^k}$, more precisely over the $I_{L^k}(\mathfrak{A})$. To make this precise we introduce some ad-hoc conventions. Predicates of arity k are naturally representable over A^k as unary predicates. For predicates of arity $1 \leqslant r < k$ we adopt a representation via the passage from $R \subseteq A^r$ to $R' := \{(\underbrace{a_1, \ldots, a_1}_{k-r}, a_1, \ldots, a_r) \mid Ra_1 \ldots a_r\} \subseteq A^k$.

Since R is first-order definable from R' and vice versa, we may in particular restrict second-order parameters Z_i as in φ above to arity k rather than $r_i \leqslant k$. The same convention is applied to global relations. In our considerations about $L^k_{\infty\omega}$-definable queries we may here restrict attention to global relations of arity k. Boolean queries can be represented by these as well if we identify 0 with $\emptyset \subseteq A^k$ and 1 with the full predicate A^k. This translation, too, is sound up to first-order interdefinability.

For $\mathfrak{A} \in \text{fin}[\tau]$, call $R \subseteq A^k$ *L^k-admissible* if it is a union of $\equiv^{L^k}$-classes over A^k. Thus by Lemma 1.33, $R \subseteq A^k$ is L^k-admissible if R is $L^k_{\infty\omega}$- (and hence also $L^k_{\omega\omega}$-) definable over $\mathfrak{A}$. Note that we are here talking about definability over an individual structure, not about definability of global relations. Any L^k-admissible $R \subseteq A^k$ is faithfully represented over $A^k/\equiv^{L^k}$ by a unary predicate

$$\underline{R} = \{\alpha \in A^k/\equiv^{L^k} \mid \overline{a} \in \alpha \Rightarrow \overline{a} \in R\}.$$

With this translation for admissible interpretations of free second-order variables we obtain a uniform translation from L^k-formulae over fin$[\tau]$ to $L^2_{\omega\omega}$-formulae over the I_{L^k} as follows. Recall that if $\varphi(Z_1,\ldots,Z_l,\overline{x})$ is in free variables Z_i and $\overline{x}$ as indicated we write $\varphi[\mathfrak{A},W_1,\ldots,W_l]$ for the predicate defined by φ in variables $\overline{x}$ over $\mathfrak{A}$ if the Z_i are interpreted by predicates W_i over $\mathfrak{A}$.

Lemma 3.16. *Let $\varphi(Z_1,\ldots,Z_l,x_1,\ldots,x_k) \in L^k_{\omega\omega}[\tau]$ with second-order variables Z_i of arity k. Then there is an $L^2_{\omega\omega}$-formula $\underline{\varphi}(\underline{Z}_1,\ldots,\underline{Z}_l,x)$ in the language of the I_{L^k} and with unary second-order variables $\underline{Z}_i$ that uniformly captures φ over the I_{L^k} in the following sense. For all $\mathfrak{A} \in$ fin$[\tau]$ and all L^k-admissible interpretations W_i for the Z_i over $\mathfrak{A}$:*

$$\underline{\varphi[(\mathfrak{A},W_1,\ldots,W_l]} = \underline{\varphi}[I_{L^k}(\mathfrak{A}),\underline{W}_1,\ldots,\underline{W}_l].$$

Proof. The proof is a straightforward induction over formulae φ. Assume without loss of generality that there is just one second-order variable Z and that each Z-atom in φ is in a tuple of mutually distinct variables (otherwise pass for instance from $Zx_1x_1\ldots$ to $\exists x_2(x_2 = x_1 \wedge Zx_1x_2\ldots)$).
i) Consider atomic φ. Let φ be a Z-atom of the admitted kind. Then $\varphi = Zx_{\rho(1)}\ldots x_{\rho(k)}$ for some $\rho \in S_k$. The formula $\underline{\varphi}(\underline{Z},x) = \exists y(S_\rho xy \wedge \underline{Z}y)$ is as desired. If φ is an atom not involving Z then it is equivalent with a finite disjunction over atomic τ-types. These translate into a disjunction over formulae $P_\theta x$ for the corresponding $\theta \in \mathrm{Atp}(\tau;k)$.
ii) Boolean operations carry over trivially.
iii) It remains to consider existential quantification. Let $\varphi = \exists x_j\psi$ and assume that $\underline{\psi}(\underline{Z},x)$ is as desired for ψ. Let $\underline{\psi}(\underline{Z},y)$ be the result of exchanging x and y throughout $\underline{\psi}$. Then the formula $\underline{\varphi}(\underline{Z},x) = \exists y(E_jxy \wedge \underline{\psi}(\underline{Z},y))$ is an adequate translation of φ. □

It follows immediately that fixed-point processes over fin$[\tau]$ translate into corresponding fixed-point processes over the I_{L^k}.

Lemma 3.17. *Let $\varphi(Z_1,\ldots,Z_l,\overline{x}) \in$ PFP$[\tau]$. Then for sufficiently large k there is a* PFP*-formula $\underline{\varphi}(\underline{Z}_1,\ldots,\underline{Z}_l,x)$ in the language of the I_{L^k} and with unary second-order variables $\underline{Z}_i$ that uniformly captures φ over the I_{L^k}. For all $\mathfrak{A} \in$ fin$[\tau]$ and all admissible interpretations W_i for the Z_i over $\mathfrak{A}$:*

$$\underline{\varphi[(\mathfrak{A},W_1,\ldots,W_l]} = \underline{\varphi}[I_{L^k}(\mathfrak{A}),\underline{W}_1,\ldots,\underline{W}_l].$$

The same holds of FP *in place of* PFP.

Sketch of Proof. The proof is obvious on the basis of the last lemma. Formally it is by induction on PFP-formulae. The PFP-step is as follows. Assume $\varphi = [\mathrm{PFP}_{X,\overline{x}}\psi]\overline{x}$ and disregard for convenience second-order parameters.

By the inductive hypothesis there is a PFP-formula $\underline{\psi}(\underline{X}, x)$ such that for appropriate k: $\underline{\psi[\mathfrak{A}, X]} = \underline{\psi}[I_{L^k}(\mathfrak{A}), \underline{X}]$ for all L^k-admissible X.

It follows inductively that the stages X_i in any fixed-point generation based on ψ are L^k-admissible predicates: the empty predicate is L^k-admissible and $\psi[\mathfrak{A}, X]$ is L^k-admissible for L^k-admissible X as it admits a representation over $A^k/\equiv^{L^k}$ through $\underline{\psi}[I_{L^k}(\mathfrak{A}), \underline{X}]$.

Let the $\underline{X}_i$ be the representations of the X_i. It is obvious that the $\underline{X}_i$ are the stages of a partial fixed-point process over I_{L^k} that is induced by $\underline{\psi}$. It follows that the partial fixed-point of the stages $\underline{X}_i$ over I_{L^k} is the representation of $\mathrm{PFP}_{X,\overline{x}}\psi$: $\underline{[\mathrm{PFP}_{X,\overline{x}}\psi]} = [\mathrm{PFP}_{\underline{X},x}\underline{\psi}]$. □

The C^k-invariants behave much like the L^k-invariants. All the information expressed in the L^k-invariant about $A^k/\equiv^{L^k}$ is expressed by the C^k-invariant about the finer representation $A^k/\equiv^{C^k}$. It is immediate therefore that the statement of the last lemma carries over to I_{C^k} in place of I_{L^k}. We state it without (C^k-admissible) second-order parameters, merely for notational convenience.

Corollary 3.18. *Let $\varphi \in \mathrm{PFP}[\tau]$, respectively $\mathrm{FP}[\tau]$. Then for sufficiently large k there is a PFP-formula, respectively FP-formula $\underline{\varphi}(x)$ in the language of the I_{C^k} that uniformly captures φ over the I_{C^k} in the sense that for all $\mathfrak{A} \in \mathrm{fin}[\tau]$: $\underline{\varphi[\mathfrak{A}]} = \underline{\varphi}[I_{C^k}(\mathfrak{A})]$.*

I_{C^k} contains numerical information encoded in the weight functions ν_j. In an extension of the statement of the last corollary we thus get that for instance the Härtig quantifier can also be captured. This will become useful later. Recall that the Härtig quantifier Q_{H} expresses cardinality equality of two definable unary predicates, cf. Definition 1.53.

Lemma 3.19. *Let $\varphi \in \mathrm{PFP}(Q_{\mathrm{H}})[\tau]$. Then for sufficiently large k there is a PFP-formula $\underline{\varphi}(x)$ in the language of the I_{C^k} that captures φ over the I_{C^k} in the sense that for all $\mathfrak{A} \in \mathrm{fin}[\tau]$: $\underline{\varphi[\mathfrak{A}]} = \underline{\varphi}[I_{C^k}(\mathfrak{A})]$.*[2]

The same holds of $\mathrm{FP}(Q_{\mathrm{H}})$ and FP in place of $\mathrm{PFP}(Q_{\mathrm{H}})$ and PFP.

Sketch of Proof. Above the proofs of Lemma 3.16 and 3.17 we only need to show that an application of the Härtig quantifier carries over to the representation on I_{C^k}. Let $\varphi = Q_{\mathrm{H}}\big((x_j; \psi_1); (x_{j'}; \psi_2)\big)$. Semantically this formula says that $\big|\{x_j | \psi_1\}\big| = \big|\{x_{j'} | \psi_2\}\big|$. Assume that there are PFP-formulae $\underline{\psi_i}(x)$ satisfying the claim of the lemma for appropriate k. Then

$$\big|\{x_j | \psi\}\big| = \textstyle\sum_{x \in \underline{\psi}} \nu_j(x)$$

is a number whose value is PTIME computable from the unary predicate $\underline{\psi}$ over the ordered domain of I_{C^k}. Hence this value is fixed-point definable over I_{C^k}, and so is equality of two such values. □

[2] Compare Definition 3.3 and Remark 3.4. It is essential that we consider the relational encoding of the full invariants, the weights ν_j inclusive.

3.4.2 The Abiteboul-Vianu Theorem

For the L^k we can already demonstrate the power of the invariants in the analysis of FP and PFP in relation to computational complexity. This leads to a theorem of Abiteboul and Vianu which is one of the major results in the field.

Definition 3.20. *Let* PTIME(I_{L^k}), *respectively* PSPACE(I_{L^k}), *stand for the class of all queries that are* PTIME, *respectively* PSPACE *computable in terms of the* I_{L^k}. *More precisely for instance for* PTIME(I_{L^k}):

(i) a boolean query Q *on* fin$[\tau]$ *is in* PTIME(I_{L^k}) *if membership of* $\mathfrak{A}$ *in* Q *is a* PTIME *property of* $I_{L^k}(\mathfrak{A})$.
(ii) an r*-ary query* R *on* fin$[\tau]$ *for* $r \leqslant k$ *is in* PTIME(I_{L^k}) *if membership of* $\overline{a}$ *in* $R^{\mathfrak{A}}$ *is a* PTIME *property of* $I_{L^k}(\mathfrak{A}, \overline{a})$. *Here* I_{L^k} *stands for the extension to an invariant on* fin$[\tau; r]$.

Equivalently, a query is in PTIME(I_{L^k}) or PSPACE(I_{L^k}) if it is $L^k_{\infty\omega}$-definable (and therefore its values $R^{\mathfrak{A}}$ will in particular be L^k-admissible over $\mathfrak{A}$) and if the natural representations $\underline{R}^{\mathfrak{A}}$ of $R^{\mathfrak{A}}$ over $A^k/\equiv^{L^k}$ are PTIME or PSPACE computable over the $I_{L^k}(\mathfrak{A})$.

Note that these classes are recursively presentable. The PTIME or PSPACE algorithms featuring in the definition are not subject to any semantic constraints: unlike the original input structures $\mathfrak{A}$, the $I_{L^k}(\mathfrak{A})$ are objects with standard encodings.

PTIME(I_{L^k}) and PSPACE(I_{L^k}) are natural classes under the following view. Consider the case of boolean queries $Q \subseteq$ fin$[\tau]$. We identify Q with its characteristic functor χ: fin$[\tau] \rightarrow \{0, 1\}$ which is subject to the condition of invariance under isomorphism. Q is $L^k_{\infty\omega}$-definable if and only if χ is in fact $\equiv^{L^k}$-invariant. This is equivalent with the existence of a presentation of χ as $\chi = \chi^* \circ I_{L^k}$ for a boolean valued mapping χ^*. Note that χ^* is defined on a set of objects with standard encodings and is not subject to any additional constraints. The same considerations apply to k-ary queries, which we may identify with isomorphism invariant boolean functors on fin$[\tau; k]$. PTIME(I_{L^k}) and PSPACE(I_{L^k}) consist exactly of those queries which are presentable by $\chi = \chi^* \circ I_{L^k}$ with PTIME or PSPACE computable functions χ^*.

It follows from the Theorems of Immerman and Vardi (Theorem 1.24) and of Abiteboul, Vardi and Vianu (Theorem 1.25) that these classes are semantically equivalent with logical systems based on FP and PFP:

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

The logics on the right-hand side consist of those formulae that are obtained as FP- or PFP-formulae applied to the FP-definable interpretations of I_{L^k} as a quotient over the k-th power. Using the fact that I_{L^k} itself is

FP-interpretable and the closure properties of FP and PFP with respect to interpretations, Lemma 1.49, $\mathrm{FP}(I_{L^k})$ and $\mathrm{PFP}(I_{L^k})$ are seen to be fragments of FP and PFP, respectively. They could obviously be characterized in purely syntactic terms if one so wishes.

Lemma 3.21. *The following semantic equivalences hold on the class of all finite relational structures:*

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

Proof. The inclusions $\mathrm{FP}(I_{L^k}) \subseteq \mathrm{FP}$ and $\mathrm{PFP}(I_{L^k}) \subseteq \mathrm{PFP}$ follow from the closure of FP and PFP under definable interpretations, Lemma 1.49. For FP-interpretability of $I_{L^k}(\mathfrak{A})$ over $\mathfrak{A}$ see Proposition 3.12. The converse inclusions $\mathrm{FP} \subseteq \bigcup_k \mathrm{FP}(I_{L^k})$ and $\mathrm{PFP} \subseteq \bigcup_k \mathrm{PFP}(I_{L^k})$ follow from Lemma 3.17. □

As a corollary to these equivalences we finally obtain the following.

Theorem 3.22 (Abiteboul, Vianu). FP $\equiv$ PFP *on the class of all finite relational structures if and only if* PTIME = PSPACE.

Proof. FP $\equiv$ PFP $\Rightarrow$ PTIME = PSPACE follows by considering the class of ordered structures and applying the theorems of Immerman, Vardi and of Abiteboul, Vardi, Vianu that equate FP with PTIME and PFP with PSPACE over these.

The real content of the theorem is the converse: if PTIME = PSPACE then FP $\equiv$ PFP over the class of *all* finite relational structures. Lemma 3.21 yields the necessary reduction of the general case to the ordered case. If PTIME = PSPACE then $\mathrm{FP} \equiv \bigcup_k \mathrm{PTIME}(I_{L^k}) \equiv \bigcup_k \mathrm{PSPACE}(I_{L^k}) \equiv \mathrm{PFP}$. □

3.4.3 Comparison of I_{C^k} and I_{L^k}

There is an obvious formal difference between the L^k- and the C^k-invariants. $I_{L^k}(\mathfrak{A})$ is interpretable as a purely relational structure over the given structures $\mathfrak{A}$. For $I_{C^k}(\mathfrak{A})$ this applies only to the relational part to which weight functions have to be added to obtain an invariant that characterizes up to $\equiv^{C^k}$. A complete relational representation of $I_{C^k}(\mathfrak{A})$ has size $|A|$, the same as $\mathfrak{A}$ itself. Setting aside our particular encoding convention, its size is at least polynomially related to the size of the original structure. The size of $I_{L^k}(\mathfrak{A})$ on the other hand is $|\,A^k/\equiv^{L^k}|$. Below, an example is reviewed of a theory in $L^3_{\infty\omega}$ which forces the size of $I_{L^k}(\mathfrak{A})$ to be logarithmically small in terms of $|A|$ in all its finite models. We have seen in the case of $L^k_{\infty\omega}$ in the last sections that the size of the invariants is directly related to the expressive power of FP. Lemma 3.21 implies that FP-evaluations over $\mathfrak{A}$ close within polynomially many steps — not in the size of $\mathfrak{A}$ but in the size of $I_{L^k}(\mathfrak{A})$

for some sufficiently large k. A similar dependence of fixed-point logic with counting on the size of the I_{C^k} will be derived in the next chapter. That this distinction is a logical phenomenon (and not just an artifact of the particular formalizations of the invariants) follows even in the very trivial case of pure sets, i.e. for $\tau = \emptyset$. Note that in this case the size of I_{L^k} is bounded by a constant, namely the number of equality types in k variables. Correspondingly, FP and all of $L^{\omega}_{\infty\omega}$ collapse to first-order logic over pure sets, see Corollary 1.32. The $C^k_{\infty\omega}$ on the other hand define arbitrarily complex classes of pure sets, and any reasonable formalization of *fixed-point logic with counting* has to render definable all PTIME arithmetical properties of the size of pure sets.

We review Example 1.16 concerning $L^3_{\infty\omega}$-definability of the class of full binary trees. It serves to show that even three variables suffice to force a logarithmic collapse in the size of I_{L^k}. The example is presented and discussed in this context by Dawar, Lindell and Weinstein in [DLW95]. The formalization in just three variables indicated in Example 1.16 is interesting because we shall see in the last chapter that no $L^2_{\infty\omega}$-theory can force a similar collapse: $k = 3$ in fact delineates the border-line for this phenomenon.

Example 3.23. By Example 1.16 there is a sentence φ in $L^3_{\infty\omega}[E]$ defining the class of full finite binary trees. Obviously the size of full binary trees is exponential in their height. The number of $L^k_{\infty\omega}$-types, however, is bounded by a polynomial in the height, since even the number of isomorphism types of k-tuples is bounded by a polynomial. The isomorphism type of a k-tuple $(v_1, \dots, v_k)$ within a given full binary tree is completely characterized by the heights of the vertices v_{ij}, $1 \leqslant i \leqslant j \leqslant k$, where v_{ij} is that vertex in which the paths from v_i and v_j to the root meet. It follows that the number of $L^k_{\infty\omega}$-types in models $\mathfrak{A} \models \varphi$ and therefore the sizes of all $I_{L^k}(\mathfrak{A})$ are polylogarithmic in the size of these models. For suitable polynomials p_k: $|I_{L^k}(\mathfrak{A})| = \left|\mathrm{Tp}^{L^k}(\mathfrak{A}; k)\right| \leqslant p_k\big(\log(|A|)\big)$ for all $\mathfrak{A} \models \varphi$.

Dawar, Lindell and Weinstein also employ tree structures with this logarithmic collapse in a padding argument to prove the second main result of Abiteboul and Vianu about the relationship between FP and PFP stated below. For the proof we refer to [DLW95]. The statement of this result is important here because we shall find the opposite for the counting extensions — the reason for this fundamental difference is that the I_{L^k} may collapse the size of structures while the I_{C^k} do not.

Let $\mathrm{PFP}|_{\mathrm{poly}}$ stand for the subclass of PFP which admits PFP-applications only where the limit is reached in a polynomially bounded number of steps. In particular $\mathrm{FP} \subseteq \mathrm{PFP}|_{\mathrm{poly}}$. Intuitively FP captures PTIME relational recursion, PFP captures PSPACE relational recursion. It would be tempting therefore to conjecture that $\mathrm{PFP}|_{\mathrm{poly}} = \mathrm{FP}$.

Theorem 3.24 (Abiteboul, Vianu).
If $\mathrm{PFP}|_{\mathrm{poly}} \subseteq \mathrm{FP}$ *then* PTIME = PSPACE.

Note that the converse implication holds as a consequence of the first theorem of Abiteboul-Vianu, Theorem 3.22.

3.5 A Partial Reduction to Two Variables

The invariants I_{C^k} and I_{L^k} have as their backbones pre-orderings defined as the stable colourings of certain graphs interpretable over the k-th power of the given structures. In the standard setting these pre-orderings themselves can be defined in $C^2_{\infty\omega}$ and $L^2_{\infty\omega}$, respectively, as shown in Section 2.2. Pursuing this connection further one can show that I_{C^k} and I_{L^k} are in fact FP-interpretable over the two-variable invariants of the game k-graphs of the underlying structures. These results later play a rôle in potential reductions for canonization problems.

Recall from Definition 2.26 that the game k-graph $\mathfrak{A}^{(k)}$ associated with $\mathfrak{A} \in \text{fin}[\tau]$ is the structure with universe A^k and with binary predicates E_j for the accessibility relations in each component and unary predicates P_θ for the identification of atomic types $\theta \in \text{Atp}(\tau; k)$. The vocabulary of $\mathfrak{A}^{(k)}$ is denoted $\tau^{(k)}$. For the technical notion of interpretability of functors compare the remarks made in connection with Example 1.47.

Proposition 3.25. *$I_{L^k}(\mathfrak{A})$ is uniformly* FP*-interpretable in $I_{L^2}(\mathfrak{A}^{(k)})$. $I_{C^k}(\mathfrak{A})$ is uniformly* FP*-interpretable in $I_{C^2}(\mathfrak{A}^{(k)})$.*

The mere functional dependencies expressed in these interpretability statements imply in particular that

$$\begin{aligned} \mathfrak{A}^{(k)} \equiv^{C^2} \mathfrak{A}'^{(k)} &\implies \mathfrak{A} \equiv^{C^k} \mathfrak{A}', \\ \mathfrak{A}^{(k)} \equiv^{L^2} \mathfrak{A}'^{(k)} &\implies \mathfrak{A} \equiv^{L^k} \mathfrak{A}'. \end{aligned}$$

The claim of the proposition goes beyond these implications, since it requires FP-interpretability or PTIME computability of one invariant in terms of the other.

Sketch of Proof. The proof is somewhat technical though not difficult. We indicate the proof for the interpretability of $I_{C^k}(\mathfrak{A})$ in $I_{C^2}(\mathfrak{A}^{(k)})$. Since we are dealing with ordered structures it suffices to show that $I_{C^k}(\mathfrak{A})$ is PTIME computable from $I_{C^2}(\mathfrak{A}^{(k)})$, compare Example 1.47.

Consider the generation of the $\preccurlyeq_i$ with limit $\preccurlyeq$, where $\preccurlyeq$ is the quotient interpretation over A^k of the ordering $(A^k/\equiv^{C^k}, \leqslant)$ underlying $I_{C^k}(\mathfrak{A})$. For each i let $(A^k/\approx_i, \leqslant_i)$ be the ordered quotient induced by $\preccurlyeq_i$.

We first show inductively how $(A^k/\approx_i, \leqslant_i)$ is interpretable in the relational part of $I_{C^2}(\mathfrak{A}^{(k)})$. For this interpretation we use those elements of $I_{C^2}(\mathfrak{A}^{(k)})$ that represent types of singletons over $\mathfrak{A}^{(k)}$, i.e. that have $x_1 = x_2$ in their atomic $\tau^{(k)}$-type. We denote this subset of the universe of $I_{C^2}(\mathfrak{A}^{(k)})$ by Δ and identify it with $\text{Tp}^{C^2}(\mathfrak{A}^{(k)}; 1)$. Recall that singletons over $\mathfrak{A}^{(k)}$ are

k-tuples over $\mathfrak{A}$. The desired interpretation is such that the $\approx_i$-class of $\overline{a} \in A^k$ is represented by the set of all $\mathrm{tp}^{C^2}_{\mathfrak{A}^{(k)}}(\overline{b})$ for $\overline{b} \approx_i \overline{a}$. It can be described by the mapping

$$\begin{array}{rcl} A^k/\approx_i & \longrightarrow & \mathcal{P}(\Delta) \\ \alpha & \longmapsto & \underline{\alpha} := \{\mathrm{tp}^{C^2}_{\mathfrak{A}^{(k)}}(\overline{b}) \mid \overline{b} \in \alpha\}, \end{array}$$

where $\mathcal{P}(\Delta)$ stands for the power set of Δ.

$(A^k/\approx_0, \leqslant_0) = (\mathrm{Atp}(\mathfrak{A};k), \leqslant_0)$ is quantifier free interpretable over Δ since each atomic type $\theta \in \mathrm{Atp}(\mathfrak{A};k)$ corresponds to those elements of Δ whose atomic $\tau^{(k)}$-type contains $P_\theta x_1$:

$$\underline{\theta} = \{\beta \in \mathrm{Tp}^{C^2}(\mathfrak{A}^{(k)};1) \mid \beta \models P_\theta x_1\}.$$

Now for the inductive step from $(A^k/\approx_i, \leqslant_i)$ to $(A^k/\approx_{i+1}, \leqslant_{i+1})$. Recall that $\preccurlyeq_{i+1}$ is determined in terms of the numbers

$$\nu_j^\alpha(\overline{a}) = \left|\{b \in A \mid \overline{a}\tfrac{b}{j} \in \alpha\}\right|$$

for $\overline{a} \in A^k$ and $\alpha \in A^k/\approx_i$. By the inductive hypothesis any such α is interpreted by a subset $\underline{\alpha}$ of Δ in $I_{C^2}(\mathfrak{A}^{(k)})$. Obviously $\nu_j^\alpha(\overline{a})$ is represented over $\mathfrak{A}^{(k)}$ as

$$\nu_j^\alpha(\overline{a}) = \left|\{\overline{b} \in A^k \mid E_j\overline{a}\overline{b} \wedge \overline{b} \in \alpha\}\right|.$$

This shows that these numbers can only depend on $\mathrm{tp}^{C^2}_{\mathfrak{A}^{(k)}}(\overline{a})$ and therefore are directly computable on $I_{C^2}(\mathfrak{A}^{(k)})$ from $\mathrm{tp}^{C^2}_{\mathfrak{A}^{(k)}}(\overline{a})$ and $\underline{\alpha}$. Thus the desired interpretation of $(A^k/\approx_{i+1}, \leqslant_{i+1})$ over $I_{C^2}(\mathfrak{A}^{(k)})$ is PTIME computable from that of $(A^k/\approx_i, \leqslant_i)$.

This refinement process terminates after polynomially many many steps and its limit is the interpretation of the ordered quotient $(A^k/\equiv^{C^k}, \leqslant)$ needed for $I_{C^k}(\mathfrak{A})$. The other data in $I_{C^k}(\mathfrak{A})$ are easily definable and computable in terms of this interpretation as follows.

The P_θ are trivially represented by atomic formulae over $\mathfrak{A}^{(k)}$.

The E_j are also atomically represented in $\mathfrak{A}^{(k)}$ and can be transferred to the interpreted $I_{C^k}(\mathfrak{A})$ as follows: α and α' are E_j-related in $I_{C^k}(\mathfrak{A})$ if they are realized by some $\overline{a}$ and $\overline{a}'$ that are E_j-related in $\mathfrak{A}^{(k)}$. Therefore $(\alpha, \alpha') \in E_j$ in $I_{C^k}(\mathfrak{A})$ if there is some

$$\beta \in \mathrm{Tp}^{C^2}(\mathfrak{A}^{(k)};2) \text{ such that } \beta \models E_j x_1 x_2\,,\ \beta|x_1 \in \underline{\alpha} \text{ and } \beta|x_2 \in \underline{\alpha}'.$$

Here $\beta|x_i$ denotes the restriction of the 2-type β to the i-th component, which is an element of $\Delta = \mathrm{Tp}^{C^2}(\mathfrak{A}^{(k)};1)$.

The weights ν_j of $I_{C^k}(\mathfrak{A})$ reduce to numerical data that are available on $I_{C^2}(\mathfrak{A}^{(k)})$ in the manner exhibited for the $\nu_j^\alpha(\overline{a})$ above.

The S_ρ are PTIME computable from the remaining data anyway according to Remark 3.5. □

Sources, attributions and remarks. As pointed out above the important concept of an ordered invariant is due to Abiteboul and Vianu [AV91]. Their invariants were abstracted from a model of relational computation and employed in an analysis FP, $\mathrm{PFP}|_{\mathrm{poly}}$ and PFP over arbitrary relational structures in terms of complexities of computations over ordered structures. The major results are the theorems of Abiteboul and Vianu, Theorems 3.22 and 3.24 above. The systematic formalization of this approach in terms of the $L^k_{\infty\omega}$ is due to Dawar [Daw93] and Dawar, Lindell and Weinstein [DLW95]. The extension and logical formulation for the $C^k_{\infty\omega}$ is presented in [GO93, Ott96a]. The corresponding applications to fixed-point logics with counting will form a main topic of the following chapter.