6. Canonization Problems

This is the first of two chapters dealing with canonization. In this chapter we consider canonization up to logical equivalences $\equiv^{\mathcal{L}}$, in particular for the logics $\mathcal{L} = L^k_{\infty\omega}$ and $C^k_{\infty\omega}$. We investigate the relation between PTIME canonization, PTIME inversion of the invariants, and the existence of recursive presentations and normal forms for related fragments of PTIME. It is shown for instance that PTIME invertibility for the I_{C^k} for all k would imply that FP+C captures exactly all queries that are PTIME computable and $C^\omega_{\infty\omega}$-definable. This and similar implications are of a hypothetical status, however: the problem of PTIME invertibility — and of PTIME canonization for C^k and L^k — remains open for arbitrary k. We show in this chapter that the general case essentially reduces to that for the *three* variable fragments. An explicit solution to the problem for the *two* variable fragments will be presented in the next chapter.

- Section 6.1 reviews the general notion of canonization and discusses canonization with respect to isomorphism in connection with algorithms on structures.

- In Section 6.2 PTIME canonization for $\equiv^{\mathcal{L}}$ is related to recursive presentations of fragments of PTIME.

- Section 6.3 discusses PTIME inversion of the I_{C^k} and I_{L^k} in relation to canonization and normal forms for the related fragments of PTIME. In particular we present theorems on the impact of PTIME invertibility of all I_{C^k}, respectively all I_{L^k} (in the sense of Definition 6.9), on the classes PTIME $\cap\, C^\omega_{\infty\omega}$ and PTIME $\cap\, L^\omega_{\infty\omega}$.

- The reduction of these results to the three variable fragments is presented in Section 6.4.

6.1 Canonization

For the general notion of canonization compare Definition 1.57 and related remarks in Section 1.7.1. Formally a function H provides canonization for $\sim$ if it satisfies two conditions. For all x we want $H(x) \sim x$ and whenever

$x \sim x'$ then $H(x) = H(x')$. Dealing with finite structures as basic objects and considering computable canonization with respect to an equivalence relation on some fin[τ], we require $H(\mathfrak{A})$ to be a structure with standard domain, $H(\mathfrak{A}) \in$ stan[τ]. Compare in particular Definition 1.61.

An important case is the canonization problem of combinatorial graph theory, namely the problem of *canonization of finite graphs up to isomorphism.* This is often also termed *graph normalization.* The same problem applies to any other class of finite structures, in particular to the entire classes fin[τ] for arbitrary finite relational vocabularies τ. Normalization for any fin[τ], however, reduces to graph normalization for most purposes. This is because there are natural encoding schemes mapping relational structures of an arbitrary fixed vocabulary to graphs in a way that would be compatible with normalization. With encodings by means of relativized interpretations of τ-structures in graphs, standardization of the parent structure (the graph) immediately induces a corresponding standardization of the interpreted τ-structure.

The problem of finding a standard representative up to isomorphism for relational structures is closely related with the analysis of algorithms over structures as discussed in the introduction (compare also Section 1.2). Standard models of computation require the input structure to be represented as a string over some alphabet. This is possible in a canonical way for ordered structures since these admit a trivial low complexity normalization procedure. Let $< \in \tau$ and recall that ord[τ] stands for the class of finite τ-structures that are linearly ordered by $<$. The natural canonization then is

$$\begin{aligned} H\colon \mathrm{ord}[\tau] &\longrightarrow \mathrm{stan}[\tau] \\ \mathfrak{A} = (A, <^{\mathfrak{A}}, \ldots) &\longmapsto (|A|, <^{|A|}, \ldots), \end{aligned}$$

where $(|A|, <^{|A|}, \ldots)$ is the unique structure in stan[τ] with the natural ordering, that is isomorphic with $(A, <^{\mathfrak{A}}, \ldots)$. If a priori we admit an arbitrary representation also for $\mathfrak{A} \in$ ord[τ] through an arbitrary isomorphic representative in stan[τ] we find that this functor H is computable in LOGSPACE.

It is difficult to imagine any feasible representation of the isomorphism type of finite structures for standard computational models and in particular the Turing model, that does not implicitly introduce a linear order on the domain of the given structure. In general one therefore has to admit representations of the abstract structure $\mathfrak{A} \in$ fin[τ] through arbitrary isomorphic representative in stan[τ] — or, equivalently, through the introduction of an arbitrary ordering for representational purposes. Uniqueness of the representative is given up and the notorious invariance problems have to be dealt with. Algorithms for structures have to satisfy a semantic invariance condition, since the outcome of the computation must be independent of the input representation.

This problem can be side-stepped, however, if there should be a feasible construction of unique representatives after all. In the general case this requires a feasibly computable functor

$$\begin{array}{rrcl} & H\colon \mathrm{fin}[\tau] & \longrightarrow & \mathrm{stan}[\tau] \\ & \mathfrak{A} & \longmapsto & H(\mathfrak{A}), \\ \text{satisfying} & \forall\mathfrak{A} & & H(\mathfrak{A}) \simeq \mathfrak{A}, \\ & \forall\mathfrak{A}\forall\mathfrak{A}' & & \mathfrak{A} \simeq \mathfrak{A}' \rightarrow H(\mathfrak{A}) = H(\mathfrak{A}'). \end{array}$$

It is here irrelevant whether we regard fin$[\tau]$ or stan$[\tau]$ as the domain of this functor. A functor H with these properties is a *computable canonization functor with respect to isomorphism* on fin$[\tau]$ in the sense of Definition 1.61. In the context of logics for fragments of PTIME "feasible" here means "PTIME computable". It is not known whether there is a PTIME normalization procedure for all finite relational structures, or equivalently for the class of all finite graphs. It is clear that PTIME graph normalization would immediately yield a PTIME algorithm for the *graph isomorphism problem.* The status with respect to complexity of the graph isomorphism problem, however, is a notorious open problem.

For an upper bound on the complexity of graph normalization one can at least show that it is contained in Δ_2^{pol} at the second level of the polynomial hierarchy. Δ_2^{pol} is the class of those problems that admit a PTIME solution relative to an oracle in NPTIME (and NPTIME $= \Sigma_1^{\mathrm{pol}}$).

Example 6.1. There is a graph normalization functor H in Δ_2^{pol}.

For the oracle we choose the *weak subgraph isomorphism problem.* The weak subgraph relation $\mathfrak{G}_1 \subseteq_{\mathrm{w}} \mathfrak{G}_2$ holds if the universe of $\mathfrak{G}_1$ is a subset of the universe of $\mathfrak{G}_2$ and if all edges of $\mathfrak{G}_1$ are also edges of $\mathfrak{G}_2$. Let $\mathcal{O}$ be the set of all standard encodings of pairs of graphs $(\mathfrak{G}_1, \mathfrak{G}_2)$ where $\mathfrak{G}_1$ is isomorphic with some $\mathfrak{G}_1' \subseteq_{\mathrm{w}} \mathfrak{G}_2$. Obviously $\mathcal{O}$ is in NPTIME, in fact it is NPTIME-complete.

Relative to the oracle $\mathcal{O}$ we get the following PTIME algorithm $\mathcal{A}$ for graph normalization.

On input (n, E), a graph on standard domain n, $\mathcal{A}$ successively computes edge relations $E_m \subseteq m \times m$ for $m = 1, \ldots, n$, where $E_1 = \emptyset$ and , for $m > 1$, E_m is the lexicographically maximal element of the set

$$S^m = \big\{ R \subseteq m \times m \;\big|\; E_{m-1} \subseteq R \text{ and } \big((m, R), (n, E)\big) \in \mathcal{O} \big\}.$$

The lexicographic ordering on the $R \subseteq m \times m$ is the usual one if R is identified with the sequence of values of its characteristic function $\chi_R(0,0)$, $\chi_R(0,1)$, $\ldots$, $\chi_R(m-1, m-1)$. It is easily shown inductively that the S^m are non-empty. All the (m, E_m) will actually be isomorphic with *subgraphs* of (n, E), as any addition of more edges to some R is an upward move in the lexicographic ordering. In fact (m, E) automatically is the lexicographically maximal graph of size m that is isomorphic with a subgraph of (n, E).

Therefore $H(n, E) := (n, E_n)$ is as desired. It remains to argue that the E_m can be determined in PTIME relative to $\mathcal{O}$, which is not quite obvious at first as in general S^m is of exponential size in m.

But to compute E_m from E_{m-1} it suffices to settle the values of $\chi = \chi_{E_m}$ at $(0, m-1), \ldots, (m-2, m-1)$. All other entries are in fact determined:

(i) $\chi(m-1, m-1) = 0$, by irreflexivity of E_m.
(ii) $\chi(m-1, j) = \chi(j, m-1)$ for $j < m-1$, by symmetry of E_m.
(iii) $\chi \restriction (m-1) \times (m-1) = \chi_{E_{m-1}}$,
because $E_m \supseteq E_{m-1}$ and $E_m \restriction (m-1) \times (m-1) \leqslant_{\text{lex}} E_{m-1}$.

That sequence $\chi(0, m-1), \ldots, \chi(m-2, m-1)$ that leads to the lexicographically maximal E_m can be constructed as follows.

Put $\chi^0(0, m-1) = \ldots = \chi^0(m-2, m-1) = 0$; the resulting E_m^0 equals E_{m-1} and thus is in S^m. Proceeding inductively, let χ^{j+1} be χ^j with the value at $(j, m-1)$ changed to 1 if the E_m^{j+1} that is so obtained is in S^m, and $\chi^{j+1} = \chi^j$ otherwise. Then $\chi := \chi^{m-1}$ is as desired.

If PTIME canonization up to isomorphism is unlikely to be attained, it is sensible to consider canonization with respect to rougher, and in particular logical notions of equivalence instead of isomorphism.

6.2 PTIME Canonization and Fragments of PTIME

Definition 6.2. *Let $\mathcal{L}$ be a logic, $\equiv^{\mathcal{L}}$ the induced notion of equivalence on* $\text{fin}[\tau]$ *and on* $\text{fin}[\tau; r]$. *A* PTIME *computable functor* $H: \text{fin}[\tau] \to \text{stan}[\tau]$ *provides* PTIME canonization up to $\equiv^{\mathcal{L}}$ *on* $\text{fin}[\tau]$ *or* canonization for $\mathcal{L}$ on $\text{fin}[\tau]$ *if the following are satisfied:*

$$\begin{array}{ll} \forall\mathfrak{A} & H(\mathfrak{A}) \equiv^{\mathcal{L}} \mathfrak{A}, \\ \forall\mathfrak{A}\forall\mathfrak{A}' & \mathfrak{A} \equiv^{\mathcal{L}} \mathfrak{A}' \to H(\mathfrak{A}) = H(\mathfrak{A}'). \end{array}$$

The analogous requirements are imposed on a functor $H: \text{fin}[\tau; r] \to \text{stan}[\tau; r]$ *for* PTIME *canonization on* $\text{fin}[\tau; r]$.

Canonization up to $\equiv^{\mathcal{L}}$ determines a unique standard representative within each class of $\mathcal{L}$-equivalent finite structures, respectively of finite structures with parameters. The difference between canonization for plain structures and structures with parameters is inessential for the logics under consideration, because $\equiv^{C^k}$ and $\equiv^{L^k}$ satisfy the requirements of the following lemma. Recall that $\equiv^{C^k}$ and $\equiv^{L^k}$ are in PTIME as relations on $\text{fin}[\tau; r]$ for all $r \leqslant k$, because the invariants for L^k and C^k on the $\text{fin}[\tau; r]$ are PTIME computable. Compare Corollaries 3.9 and 3.14.

Lemma 6.3. *Let $\mathcal{L}$ be such that $\equiv^{\mathcal{L}}$ is in* PTIME *as a relation on* $\text{fin}[\tau; r]$ *and such that $\mathfrak{A} \equiv^{\mathcal{L}} \mathfrak{A}'$ implies that $\mathfrak{A}$ and $\mathfrak{A}'$ realize the same $\mathcal{L}$-types of r-tuples:* $\mathfrak{A} \equiv^{\mathcal{L}} \mathfrak{A}' \Rightarrow \text{Tp}^{\mathcal{L}}(\mathfrak{A}; r) = \text{Tp}^{\mathcal{L}}(\mathfrak{A}'; r)$.

Then any PTIME *canonization functor* $H: \text{fin}[\tau] \to \text{stan}[\tau]$ *extends naturally to a* PTIME *canonization* $H: \text{fin}[\tau; r] \to \text{stan}[\tau; r]$ *on* $\text{fin}[\tau; r]$.

Sketch of Proof. The following extension of H satisfies the requirements — we denote it H as well. Let $H(\mathfrak{A},\bar{a}) := \big(H(\mathfrak{A}),\bar{b}\big)$ where $\bar{b}$ is the lexicographically least r-tuple over the standard domain of $H(\mathfrak{A})$ for which $\big(H(\mathfrak{A}),\bar{b}\big) \equiv^{\mathcal{L}} (\mathfrak{A},\bar{a})$. □

PTIME canonization bears the following simple yet fundamental relationship with recursive presentations of fragments of PTIME.

Lemma 6.4. *Let H provide* PTIME *canonization up to $\equiv^{\mathcal{L}}$ on* $\text{fin}[\tau]$ *and on the* $\text{fin}[\tau;r]$. *Then the class of all those queries over* $\text{fin}[\tau]$ *that are* PTIME *computable in the usual sense* and *closed with respect to $\equiv^{\mathcal{L}}$, is recursively enumerable. In fact the following are equivalent for any boolean query $Q \subseteq$* $\text{fin}[\tau]$*:*

(i) Q is closed with respect to $\equiv^{\mathcal{L}}$ and Q is in PTIME.
(ii) $Q = \big\{\mathfrak{A} \mid H(\mathfrak{A}) \in Q\big\}$ and there is a PTIME *algorithm that recognizes* $Q \cap \text{stan}[\tau]$.

For an r-ary global relation R on $\text{fin}[\tau]$ *the following are equivalent:*

(i) R is closed with respect to $\equiv^{\mathcal{L}}$ and PTIME*-computable.*
(ii) $R^{\mathfrak{A}} = \big\{\bar{a} \in A \mid \bar{b} \in R^{\mathfrak{B}}$ where $(\mathfrak{B},\bar{b}) = H(\mathfrak{A},\bar{a})\big\}$ and there is a PTIME *algorithm that, applied to $H(\mathfrak{A},\bar{a}) = (\mathfrak{B},\bar{b})$, decides whether $\bar{b} \in R^{\mathfrak{B}}$.*

Note that the algorithms in (ii) are not subject to any semantic constraints, since these algorithms need merely realize boolean *functions* on $\text{stan}[\tau]$ or $\text{stan}[\tau;r]$, respectively. A natural recursive set of representatives consists of all algorithms that first check the input size, then initialize some counter to a fixed polynomial in this size and terminate their computation after this pre-set number of steps (polynomially clocked algorithms).

Sketch of Proof. We indicate the proof for boolean queries. Observe that $\equiv^{\mathcal{L}}$–closure is equivalent with $Q = \big\{\mathfrak{A} \mid H(\mathfrak{A}) \in Q\big\}$. Therefore, any PTIME algorithm $\mathcal{A}$ that recognizes an $\equiv^{\mathcal{L}}$–closed class Q is semantically equivalent with $\mathcal{A} \circ H$. For (i) $\Rightarrow$ (ii) use $\mathcal{A}$ in restriction to $\text{stan}[\tau]$. For the converse use $\mathcal{A}$ as given in (ii) and compose it with H to get the PTIME algorithm $\mathcal{A} \circ H$ which computes the boolean query Q over $\text{fin}[\tau]$. □

Definition 6.5. *Let* PTIME$\cap\mathcal{L}$ *stand for the class of all those global relations that are both in* PTIME *and $\mathcal{L}$-definable.*

Recall from Lemma 1.33 that for logics $\mathcal{L}$ that are closed under countable disjunctions and conjunctions and under negation, $\mathcal{L}$-definability coincides with closure under $\equiv^{\mathcal{L}}$. Lemma 6.4 therefore yields a connection between PTIME canonization for $\mathcal{L}$ and a recursive presentation for PTIME$\cap\mathcal{L}$. Assume for the following definition that $H\colon \text{fin}[\tau] \to \text{stan}[\tau]$ provides canonization on $\text{fin}[\tau]$ and extends to functors $H\colon \text{fin}[\tau;r] \to \text{stan}[\tau;r]$ on the $\text{fin}[\tau;r]$ in the sense of Lemma 6.3 above.

Definition 6.6. *Let* PTIME(H) *stand for the class of all queries that are* PTIME *computable in terms of the images under* H*. More precisely,*

(i) a boolean query Q *on* $\mathrm{fin}[\tau]$ *is in* PTIME(H) *if membership of* $\mathfrak{A}$ *in* Q *is a* PTIME *property of* $H(\mathfrak{A})$.

(ii) an r*-ary query* R *on* $\mathrm{fin}[\tau]$ *is in* PTIME(H) *if membership of* $\overline{a}$ *in* $R^{\mathfrak{A}}$ *is a* PTIME *property of* $H(\mathfrak{A}, \overline{a})$.

Again, as the PTIME algorithms mentioned in the definition are not subject to additional semantic constraints, PTIME(H) is recursively presented through all compositions of polynomially clocked algorithms with some fixed algorithm for H.

Lemma 6.4 can be rephrased with this notion of PTIME(H) as follows. We state it for $\mathcal{L} = C^k_{\infty\omega}$ or $L^k_{\infty\omega}$. Note that for these any PTIME canonization on $\mathrm{fin}[\tau]$ extends to all $\mathrm{fin}[\tau; r]$ with $r \leqslant k$ by Lemma 6.3. This is sufficient for the statement below since there are no $L^k_{\infty\omega}$- or $C^k_{\infty\omega}$-definable queries in arities greater than k.

Corollary 6.7. *Let* $\mathcal{L} = C^k_{\infty\omega}$ *or* $L^k_{\infty\omega}$ *and let* H *provide* PTIME *canonization for* $\mathcal{L}$ *on* $\mathrm{fin}[\tau]$. *Then* H *extends to the* $\mathrm{fin}[\tau; r]$ *for* $r \leqslant k$ *and*

$$\text{PTIME} \cap \mathcal{L} \equiv \text{PTIME}(H).$$

In particular PTIME $\cap\, \mathcal{L}$ *is recursively enumerable (i.e. admits a recursive presentation).*

6.3 Canonization and Inversion of the Invariants

As sketched in the abstract setting in Lemma 1.60, canonization problems are generally related with inversion problems for complete invariants. While a canonization H must assign representatives, complete invariants may assign any kind of values that are characteristic of classes. Canonization may be obtained from an invariant if it is possible to reconstruct a typical member of each class on the basis of the value of that class under the invariant. The mere existence of such an inverse is obvious from the definitions. Its complexity, however, is critical. Different complete invariants for the same equivalence relation might lead to entirely different inversion problems in particular with respect to complexity.

We return to the canonization problem for the $C^k_{\infty\omega}$ and $L^k_{\infty\omega}$. The functors I_{C^k} and I_{L^k} provide complete invariants. Recall that we write I_{C^k} and I_{L^k} for the complete invariants on $\mathrm{fin}[\tau]$ as well as for their natural extensions to the $\mathrm{fin}[\tau; r]$ for $r \leqslant k$. We shall see below that also with respect to the corresponding inversion problems a solution for I_{L^k} or I_{C^k} on $\mathrm{fin}[\tau]$ naturally extends to a solution over the $\mathrm{fin}[\tau; r]$. We restate for convenience the definition of an inverse to a complete invariant, Definition 1.59, in the present context.

Definition 6.8. *Let $\mathcal{L} = C^k_{\infty\omega}$ or $L^k_{\infty\omega}$, $I_{\mathcal{L}}$ the corresponding invariant. A function*

$$F: \{I_{\mathcal{L}}(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{fin}[\tau]\} \longrightarrow \mathrm{stan}[\tau]$$

is an inverse for $I_{\mathcal{L}}$ *on* fin[τ] *if it satisfies:* $\forall \mathfrak{A}\ F(I_{\mathcal{L}}(\mathfrak{A})) \equiv^{\mathcal{L}} \mathfrak{A}$. *The analogous condition applies for inverses of $I_{\mathcal{L}}$ on* fin[$\tau; r$] *for $r \leqslant k$. Equivalently these conditions can be put as $I_{\mathcal{L}} \circ F = id$ on image($I_{\mathcal{L}}$).*

Generally an inversion of a complete invariant yields a canonization simply through composition of the inverse with the invariant itself (Lemma 1.60). Also a PTIME computable inversion F here yields PTIME canonization, since the $I_{\mathcal{L}}$ themselves are PTIME computable. Note however that the converse need not a priori be true. It is conceivable that H provides PTIME canonization while the associated F defined by the requirement that $H = F \circ I_{\mathcal{L}}$ might not be in PTIME. In fact, for the $L^k_{\infty\omega}$ with $k \geqslant 3$ we already know that inversion of I_{L^k} cannot be in PTIME in the usual sense, simply because the image under F might necessarily be of a size that is exponential in the size of the argument. See Example 3.23. The following definition takes care of this obvious obstacle and defines PTIME inversion for the I_{L^k} as an inversion that is polynomial time computable in terms of the size of the desired image.

Definition 6.9. *We say that I_{C^k}* admits PTIME inversion *if there is an inverse F for I_{C^k} that is* PTIME *computable in the usual sense.*

I_{L^k} admits PTIME inversion *if there is an inverse F for I_{L^k}, such that for all $\mathfrak{A}$, F is computable on $I_{L^k}(\mathfrak{A})$ in time polynomial in* $\min\{|\mathfrak{B}| \mid \mathfrak{B} \equiv^{L^k} \mathfrak{A}\}$. *To mark the difference in the complexity requirement let us say that such F is computable in* PTIME*.

In either case we shall speak, however, of F as a PTIME inverse of the invariant.

PTIME canonization and PTIME inversion for the L^k are discussed in Dawar's dissertation [Daw93] and in [DLW95]. The appropriate notion of PTIME inversion of I_{L^k} is put forward there and the question whether I_{L^k} admits PTIME inversion in this sense is formulated as an open problem.

A natural and intuitively stronger definition of PTIME inversion for I_{L^k} would be to require an algorithm that takes as its input pairs $(I_{L^k}(\mathfrak{A}), n)$ and produces in time polynomial in $\max\{|I_{L^k}(\mathfrak{A})|, n\}$ a structure $\mathfrak{B} \in \mathrm{stan}[\tau]$ of size n with $\mathfrak{B} \equiv^{L^k} \mathfrak{A}$ if such exists. From such an algorithm a PTIME inverse in the sense of the preceding definition is obtained through application to $(I_{L^k}(\mathfrak{A}), n)$ for growing n until a successful output is constructed. This exhaustive search for a standardized pre-image under I_{L^k} (of minimal size even) is still polynomial in the size of a minimal solution.

We come to the extension of inverses to the I_{L^k} and I_{C^k} on fin[τ] to inverses of the extended invariants on fin[$\tau; r$], $r \leqslant k$.

Lemma 6.10. *Let* $\mathcal{L} = C^k_{\infty\omega}$ *or* $L^k_{\infty\omega}$ *and* $I_{\mathcal{L}} = I_{C^k}$ *or* I_{L^k}*, respectively.* PTIME *inversion of* $I_{\mathcal{L}}$ *on* fin$[\tau]$ *extends naturally to* PTIME *inversion of* $I_{\mathcal{L}}$ *on* fin$[\tau; r]$ *for* $r \leqslant k$.

Sketch of Proof. Assume $F\colon \{I_{\mathcal{L}}(\mathfrak{A}) \mid \mathfrak{A} \in \text{fin}[\tau]\} \to \text{stan}[\tau]$ is an inverse to $I_{\mathcal{L}}$ on fin$[\tau]$. Recall that the extension of $I_{\mathcal{L}}$ to fin$[\tau; r]$ maps $(\mathfrak{A}, \overline{a})$ to the expansion of $I_{\mathcal{L}}(\mathfrak{A})$ in which the $\mathcal{L}$-type of $\overline{a}$ is marked: $I_{\mathcal{L}}(\mathfrak{A}, \overline{a}) = \big(I_{\mathcal{L}}(\mathfrak{A}), [\overline{a}]\big)$. Extend F to fin$[\tau; r]$ by putting $F\big(I_{\mathcal{L}}(\mathfrak{A}, \overline{a})\big) := \big(F(I_{\mathcal{L}}(\mathfrak{A})), \overline{b}\big)$ for the lexicographically least tuple $\overline{b}$ in the standard domain of $F(I_{\mathcal{L}}(\mathfrak{A}))$ which satisfies $I_{\mathcal{L}}\big(F(I_{\mathcal{L}}(\mathfrak{A})), \overline{b}\big) = I_{\mathcal{L}}(\mathfrak{A}, \overline{a})$. The search for this tuple is polynomially bounded in the size of $F(I_{\mathcal{L}}(\mathfrak{A}))$. Therefore if F is PTIME, respectively PTIME* in the sense of the preceding definition, on fin$[\tau]$, then so is its extension to fin$[\tau; r]$. □

Theorem 6.11. *Let* $\mathcal{L} = C^k_{\infty\omega}$ *or* $L^k_{\infty\omega}$*, and correspondingly* $I_{\mathcal{L}} = I_{C^k}$ *or* I_{L^k}*. If* F *is a* PTIME *inverse for* $I_{\mathcal{L}}$*, then* $H := F \circ I_{\mathcal{L}}$ *provides* PTIME *canonization for* $\mathcal{L}$*. Moreover this composition is compatible with the respective natural extensions of the* $I_{\mathcal{L}}$*,* H *and* F *to the* fin$[\tau; r]$ *for* $r \leqslant k$.

Sketch of Proof. We check the requirements in the case of I_{L^k}. H maps $\mathfrak{A} \in$ fin$[\tau]$ to a standard structure equivalent with $\mathfrak{A}$, since $I_{L^k} \circ F \circ I_{L^k} = I_{L^k}$ by the definition of inverses. As a composition with I_{L^k}, H certainly maps L^k-equivalent structures to the same image. It remains to check that H is in PTIME, even if F is computable only in PTIME* in the sense of Definition 6.9. Since $\min\big\{|\mathfrak{B}| \bigm| \mathfrak{B} \equiv^{L^k} \mathfrak{A}\big\} \leqslant |\mathfrak{A}|$, the computation of F on $I_{L^k}(\mathfrak{A})$ is still polynomial in terms of $|\mathfrak{A}|$.

Compatibility with the extensions to cover fin$[\tau; r]$ instead of fin$[\tau]$ follows directly from the definition of these extensions. See in particular the above lemma and compare with Lemma 6.3. □

Combining Corollary 6.7 with Theorem 6.11 we get the following connection between PTIME inversion of the I_{C^k} and the capturing of PTIME $\cap\, C^k_{\infty\omega}$. Recall Definition 4.17 for the classes PTIME$\big(I_{C^k}\big)$. The global relations in PTIME$\big(I_{C^k}\big)$ are those that are PTIME computable over the invariants I_{C^k}. Logically the same class is representable by the logics FP$\big(I_{C^k}\big)$ also discussed in connection with Definition 4.17 and Theorem 4.18.

Theorem 6.12. *If* I_{C^k} *admits* PTIME *inversion, then*

$$\text{PTIME} \cap C^k_{\infty\omega} \equiv \text{PTIME}\big(I_{C^k}\big) \equiv \text{FP}\big(I_{C^k}\big).$$

Proof. Let F be a PTIME inverse for I_{C^k}. The non-trivial inclusion PTIME $\cap$ $C^k_{\infty\omega} \subseteq \text{PTIME}\big(I_{C^k}\big)$ follows from Corollary 6.7 if we observe that $F \circ I_{C^k}$ provides PTIME canonization for $C^k_{\infty\omega}$: $\text{PTIME} \cap C^k_{\infty\omega} \subseteq \text{PTIME}\big(F \circ I_{C^k}\big) \subseteq \text{PTIME}\big(I_{C^k}\big)$, as F is in PTIME. □

Putting Theorem 4.18 — $\text{FP+C} \equiv \bigcup_k \text{FP}(I_{C^k})$ — and the last theorem together, we obtain the following hypothetical theorem.

Corollary 6.13. *If the I_{C^k} admit* PTIME *inversion for all k, then*

$$\text{PTIME} \cap C^{\omega}_{\infty\omega} \equiv \text{FP+C}.$$

For the $L^k_{\infty\omega}$ the situation is somewhat less smooth because of the possible collapse in size that can occur in the passage from $\mathfrak{A}$ to $I_{L^k}(\mathfrak{A})$. Note that, because of this potential collapse, it is not true that — as in the proof of Theorem 6.12 — $\text{PTIME}(F \circ I_{L^k}) \subseteq \text{PTIME}(I_{L^k})$ for a PTIME inverse F of I_{L^k} (which need actually only be computable in PTIME*).

Theorem 6.14. *Assume I_{L^k} admits* PTIME *inversion through F. Then* PTIME $\cap\, L^k_{\infty\omega}$ *is recursively enumerable. A boolean query is in* PTIME $\cap\, L^k_{\infty\omega}$ *if it is computable on the basis of the $I_{L^k}(\mathfrak{A})$ in time polynomial in the size of $F(I_{L^k}(\mathfrak{A}))$. Similarly for the computation of an r-ary global relation in terms of the $I_{L^k}(\mathfrak{A}, \overline{a})$.*

Let us say that I_{L^k} *is bounded* on a class $\mathcal{K} \subseteq \text{fin}[\tau]$ if there is a polynomial p such that $|\mathfrak{A}| \leqslant p(|I_{L^k}(\mathfrak{A})|)$ for all $\mathfrak{A} \in \mathcal{K}$. Obviously, if I_{L^k} is bounded on $\mathcal{K}$, then so is $I_{L^{k'}}$ for all $k' > k$. Suppose that I_{L^k} is bounded on $\mathcal{K}$ and that I_{L^k} admits PTIME inversion through F. Then F must in fact be computable in PTIME rather than in PTIME*: $F(I_{L^k}(\mathfrak{A}))$ must be polynomial time computable in terms of $|A|$ by definition, and $|A|$ is polynomial in the size of $I_{L^k}(\mathfrak{A})$ for bounded I_{L^k}. The following is then proved in precise analogy with Corollary 6.13 above.

Corollary 6.15. *If the I_{L^k} admit* PTIME *inversion for all k, then*

$$\text{PTIME} \cap L^{\omega}_{\infty\omega} \equiv \text{FP} \quad \textit{on } \mathcal{K}$$

for all classes $\mathcal{K}$ on which I_{L^k} is bounded for some k.

6.4 A Reduction to Three Variables

We exhibit a reduction technique that shows that PTIME canonization and PTIME inversion for the $L^k_{\infty\omega}$ and $C^k_{\infty\omega}$ with arbitrary k essentially reduce to the three variable cases. *'Essentially'* because the proposed reduction does not work in a k-by-k fashion but rather introduces a shift in the number of variables of the following kind. Assuming for instance PTIME invertibility of I_{C^3} we get a PTIME construction that, given $I_{C^m}(\mathfrak{A})$ for certain $m > k$, yields a standard structure that is C^k-equivalent with $\mathfrak{A}$. The effect of this mismatch is smoothed out, however, if we consider the effect with respect to the unions across all levels k. For instance, from PTIME invertibility of I_{C^3} we shall still get FP+C $\equiv$ PTIME $\cap\, C^{\omega}_{\infty\omega}$.

Here are the precise statements concerning the reduction, first in terms of canonization, then in terms of inversion of the invariants. Note that a priori these statements might be of independent interest, since the existence of a PTIME canonization procedure does not, as far as we can see, imply PTIME invertibility of the particular invariants considered here.

Theorem 6.16. *Suppose $C^3_{\infty\omega}$ admits* PTIME *canonization. Then there is for each τ and each k a* PTIME *functor $K^{3k,k}$: $\mathrm{fin}[\tau] \longrightarrow \mathrm{stan}[\tau]$ such that*

$$\begin{aligned} &\forall\mathfrak{A} \quad K^{3k,k}(\mathfrak{A}) \equiv^{C^k} \mathfrak{A}, \\ &\forall\mathfrak{A}\forall\mathfrak{A}' \quad \mathfrak{A} \equiv^{C^{3k}} \mathfrak{A}' \rightarrow K^{3k,k}(\mathfrak{A}) = K^{3k,k}(\mathfrak{A}'). \end{aligned}$$

The same statement holds of the $L^k_{\infty\omega}$ and under the assumption that $L^3_{\infty\omega}$ admits PTIME *canonization.*

Note the difference between $K^{3k,k}$ and a canonization functor with respect to $C^k_{\infty\omega}$. While $K^{3k,k}$ also produces C^k-equivalent standard structures, these representatives may depend not only on the C^k-theory of the given structure but on its C^{3k}-theory. The analogous reduction result for the inversion problem is the following.

Theorem 6.17. *Suppose I_{C^3} admits* PTIME *inversion. Then there is for each τ and each k a* PTIME *functor $G^{3k,k}$: image($I_{C^{3k}}$) $\longrightarrow$ $\mathrm{stan}[\tau]$ such that*

$$\forall\mathfrak{A} \quad G^{3k,k}\big(I_{C^{3k}}(\mathfrak{A})\big) \equiv^{C^k} \mathfrak{A},$$

or equivalently, $I_{C^3} \circ G^{3k,k} = \Pi^{3k,k}$ where $\Pi^{3k,k}$ is the obvious projection that sends $I_{C^{3k}}(\mathfrak{A})$ to $I_{C^k}(\mathfrak{A})$. Again, the same holds (for PTIME* *computability) with respect to the $L^k_{\infty\omega}$ and I_{L^k} and under the assumption that I_{L^3} admits* PTIME *inversion.*

The appropriate notion of PTIME* computability for $G^{3k,k}$: image$(I_{L^{3k}}) \rightarrow \mathrm{stan}[\tau]$ is the following: $G^{3k,k}\big(I_{L^{3k}}(\mathfrak{A})\big)$ has to be computable in time polynomial in the size of a minimal $\mathfrak{B}$ that is L^{3k}-equivalent with $\mathfrak{A}$.

For our present purposes we thus have the following corollaries. The first is in terms of PTIME canonization for the three variable case, the second in terms of PTIME inversion for the three variable invariants. In both settings we find that the general statements of Corollaries 6.7, and Corollaries 6.13 and 6.15 respectively, reduce to the three variable cases if we consider the overall effect on the unions across all k, PTIME $\cap\, C^\omega_{\infty\omega}$ and PTIME $\cap\, L^\omega_{\infty\omega}$.

Corollary 6.18. *Let $\mathcal{L} = C^\omega_{\infty\omega}$ or $L^\omega_{\infty\omega}$, respectively. Assume that $C^3_{\infty\omega}$, respectively $L^3_{\infty\omega}$, admits* PTIME *canonization. Then* PTIME$\cap\mathcal{L}$ *is recursively enumerable; in fact*

$$\mathrm{PTIME} \cap \mathcal{L} \equiv \bigcup_k \mathrm{PTIME}\big(K^{3k,k}\big),$$

where PTIME$\big(K^{3k,k}\big)$ *is formally defined in analogy with Definition 6.6 for the functors $K^{3k,k}$ as characterized in Theorem 6.16 (and their natural extensions to the* $\mathrm{fin}[\tau; r]$ *where r-ary queries rather than just boolean ones are concerned).*

Sketch of Proof. Consider $\mathcal{L} = C^{\omega}_{\infty\omega}$. Both inclusions are in fact obvious. Any query in some $\text{PTIME}(K^{3k,k})$ is in PTIME, and also in $C^{3k}_{\infty\omega}$ since the image under $K^{3k,k}$ only depends on the $C^{3k}_{\infty\omega}$-theory of structures. For the converse inclusion note that any query in $\text{PTIME} \cap C^{k}_{\infty\omega}$ can without affecting its semantics be evaluated after application of $K^{3k,k}$ because $K^{3k,k}$ preserves C^k-equivalence. □

Corollary 6.19. *(i) Assuming that I_{C^3} admits* PTIME *inversion we get:* $\text{PTIME} \cap C^{\omega}_{\infty\omega} \equiv \text{FP+C}$.

(ii) If I_{L^3} admits PTIME *inversion, then* $\text{PTIME} \cap L^{\omega}_{\infty\omega} \equiv \text{FP}$ *on all classes on which I_{L^k} is bounded for some k.*

Sketch of Proof. We indicate how the claim for boolean queries in $\text{PTIME} \cap C^{\omega}_{\infty\omega}$ follows from Theorem 6.17. Suppose that $Q \subseteq \text{fin}[\tau]$ is in $\text{PTIME} \cap C^{k}_{\infty\omega}$. Let $\mathcal{A}$ be a PTIME algorithm that recognizes Q. Then $\mathfrak{A} \in Q$ if and only if $G^{3k,k}(I_{C^{3k}}(\mathfrak{A})) \in Q$ if and only $\mathcal{A} \circ G^{3k,k} \circ I_{C^{3k}}$ accepts $\mathfrak{A}$. The latter composition is in FP+C because $I_{C^{3k}}$ is FP+C-interpretable over the $\mathfrak{A}^*$ and $\mathcal{A} \circ G^{3k,k}$ is FP-interpretable as a PTIME functor on the ordered $I_{C^{3k}}$. Closure of FP+C under interpretations (Proposition 4.8) yields $Q \in \text{FP+C}$. The converse inclusion $\text{FP+C} \subseteq \text{PTIME} \cap C^{\omega}_{\infty\omega}$ is obvious anyway (compare Corollary 4.20).

For the case of $\text{PTIME} \cap L^{\omega}_{\infty\omega}$ compare the appropriate modifications in Corollary 6.15 to adapt the argument to obtain (ii). □

Whether or not the three variable cases are solvable, remains open. The reduction achieved here therefore remains hypothetical. The two variable case is settled positively in the next chapter. In view of the above statements, a positive solution in the three variable case would be a major break-through in the understanding of the bounded-variable fragments of PTIME. The reduction argument itself is of interest because it also applies to other model theoretic questions about the $L^{k}_{\infty\omega}$ and $C^{k}_{\infty\omega}$, in particular we think of questions related to spectrum properties for these fragments, cf. [Ott96b]. For the present investigation it also illustrates where the essential power of *three*, as compared to *two* variables lies. At a more technical level it may also indicate potential obstacles for three variable canonization.

It might be worth pointing out that 3 is just the minimal number of variables for which we can show the reduction to go through. The reduction argument applies, essentially unchanged, to any other number of variables above 3. (And indeed, it is not clear why for instance k-variable canonization for some $k > 3$ should directly yield 3-variable canonization.)

None of the material in the rest of this chapter will be used in the last chapter on two-variable canonization.

6.4.1 The Proof of Theorems 6.16 and 6.17

The following definition of the *k-th power* of a relational structure resembles the definition of the game k-graphs, Definition 2.26. Here we include more complete information about the equality types of pairs of k-tuples for reasons that will become apparent below.

For finite relational τ let $\tau^{[k]}$ consist of unary predicates P_θ for $\theta \in \mathrm{Atp}(\tau; k)$ and binary predicates $\overset{i,j}{=}$ for $1 \leqslant i, j \leqslant k$. The intended interpretation for the latter — over some A^k — is that $(\overline{a}, \overline{a}') \in \overset{i,j}{=}$ if $a_i = a'_j$. We shall write $\overline{a} \overset{i,j}{=} \overline{a}'$ instead of $(\overline{a}, \overline{a}') \in \overset{i,j}{=}$.

Definition 6.20. *For $\mathfrak{A} \in \mathrm{fin}[\tau]$ let the k-*th power *of $\mathfrak{A}$ be the following structure $\mathfrak{A}^{[k]}$ in vocabulary $\tau^{[k]}$:*

$$\mathfrak{A}^{[k]} = \Big(A^k, (\overset{i,j}{=}), (P_\theta)\Big),$$

with the natural interpretations for the P_θ and the $\overset{i,j}{=}$. Denote by Γ_k the functor that takes $\mathfrak{A}$ to its k-th power $\mathfrak{A}^{[k]}$. Let $\Gamma_k\big(\mathrm{fin}[\tau]\big) \subseteq \mathrm{fin}[\tau^{[k]}]$ denote the closure under isomorphisms of the class of all k-th powers $\mathfrak{A}^{[k]}$ for $\mathfrak{A} \in \mathrm{fin}[\tau]$.

Just as the game k-graphs $\mathfrak{A}^{(k)}$, the $\mathfrak{A}^{[k]}$ are quantifier free interpretable in the k-th power over the given structures $\mathfrak{A}$.[1] Moreover, the game k-graphs $\mathfrak{A}^{(k)}$ are quantifier free (and directly) interpretable over the $\mathfrak{A}^{[k]}$: the edge relation E_j of the game k-graphs is the intersection of the $\overset{i,i}{=}$ for all $i \neq j$. Note, however, that conversely the $\overset{i,j}{=}$ are not quantifier free definable from the E_j.

The crucial fact for the desired reduction is that the C^k-theory (respectively L^k-theory) of $\mathfrak{A}$ is fully captured by the C^2-theory (respectively L^2-theory) of the k-th power $\mathfrak{A}^{[k]}$ of $\mathfrak{A}$. This follows directly from Proposition 3.25 where it was shown that even the C^2-theory of the game k-graph $\mathfrak{A}^{(k)}$ determines the C^k-theory of $\mathfrak{A}$. Clearly, the C^2-theory of $\mathfrak{A}^{[k]}$ determines that of $\mathfrak{A}^{(k)}$ owing to quantifier free interpretability of $\mathfrak{A}^{(k)}$ in $\mathfrak{A}^{[k]}$. In fact Proposition 3.25 says that the k-variable invariants $I_{C^k}(\mathfrak{A})$ or $I_{L^k}(\mathfrak{A})$ are PTIME computable (FP-interpretable) in the 2-variable invariants of the game k-graphs, $I_{C^2}\big(\mathfrak{A}^{(k)}\big)$ or $I_{L^2}\big(\mathfrak{A}^{(k)}\big)$. This carries over to the $\mathfrak{A}^{[k]}$ as well as for instance $I_{C^2}\big(\mathfrak{A}^{(k)}\big)$ is PTIME computable (FP-interpretable) in $I_{C^2}\big(\mathfrak{A}^{[k]}\big)$. We thus have the following, as a corollary to Proposition 3.25.

Proposition 6.21. *The two-variable theories of the k-th powers fully determine the k-variable theories of the base structures:*

$$\mathfrak{A}^{[k]} \equiv^{L^2} \mathfrak{A}'^{[k]} \Rightarrow \mathfrak{A} \equiv^{L^k} \mathfrak{A}', \quad \textit{and} \quad \mathfrak{A}^{[k]} \equiv^{C^2} \mathfrak{A}'^{[k]} \Rightarrow \mathfrak{A} \equiv^{C^k} \mathfrak{A}'.$$

Moreover, $I_{C^k}(\mathfrak{A})$ and $I_{L^k}(\mathfrak{A})$ are PTIME *computable from $I_{C^2}\big(\mathfrak{A}^{[k]}\big)$ and $I_{L^2}\big(\mathfrak{A}^{[k]}\big)$, respectively.*

[1] There is no conflict with the notion of *interpretability in the k-th power*: this notion may be identified with (direct) interpretability over the (interpreted) $\mathfrak{A}^{[k]}$.

Canonization or inversion of the invariants in the three variable case will prove to be sufficient for Theorems 6.16 and 6.17 because being a k-th power is definable in three variables, in fact even in $L^3_{\omega\omega}$.

Lemma 6.22. *Let the arities in τ be at most k.*

(i) There is a sentence φ in $L^3_{\omega\omega}[\tau^{[k]}]$ such that $\mathrm{fmod}(\varphi) = \Gamma_k(\mathrm{fin}[\tau])$.
(ii) $\mathfrak{A}$ is PTIME *computable from $\mathfrak{A}^{[k]}$.*

More precisely, (ii) is to say that there is a PTIME algorithm that maps $\mathfrak{C} \in \mathrm{stan}[\tau^{[k]}] \cap \Gamma_k(\mathrm{fin}[\tau])$ to a structure $\mathfrak{B} \in \mathrm{stan}[\tau]$ such that $\mathfrak{B}^{[k]} \simeq \mathfrak{C}$.

The proof of the lemma is postponed — we first show how it applies to prove Theorems 6.16 and 6.17. For this we need one more simple lemma about an interpretability relation between certain invariants.

Lemma 6.23. *Let $m \geqslant 2$. Then the m-variable theories of of the k-th powers are fully determined by the mk-variable theories of the base structures:*

$$\mathfrak{A} \equiv^{L^{mk}} \mathfrak{A}' \Rightarrow \mathfrak{A}^{[k]} \equiv^{L^m} \mathfrak{A}'^{[k]}, \quad \textit{and} \quad \mathfrak{A} \equiv^{C^{mk}} \mathfrak{A}' \Rightarrow \mathfrak{A}^{[k]} \equiv^{C^m} \mathfrak{A}'^{[k]}.$$

Moreover, $I_{C^m}(\mathfrak{A}^{[k]})$ and $I_{L^m}(\mathfrak{A}^{[k]})$ are PTIME *computable from $I_{C^{mk}}(\mathfrak{A})$ and $I_{L^{mk}}(\mathfrak{A})$, respectively.*

Proof. The proof is similar to that of Proposition 3.25: it suffices to check that the entire inductive generation of the pre-ordering underlying the m-variable invariant of $\mathfrak{A}^{[k]}$ can be simulated over the mk-variable invariant of $\mathfrak{A}$. Let $\approx_i$ and $\preccurlyeq_i$ be the stages in the generation of $\equiv^{C^m}$ and $\preccurlyeq$ over $\mathfrak{A}^{[k]}$ as required for $I_{C^m}(\mathfrak{A}^{[k]})$.

Writing $\overline{a} = (\overline{a}^{(1)}, \ldots, \overline{a}^{(m)})$ for mk-tuples over A we indicate their identification with m-tuples over A^k. A $\approx_i$-class α can be represented over $I_{C^{mk}}(\mathfrak{A})$ as

$$\underline{\alpha} = \{\mathrm{tp}^{C^{mk}}_{\mathfrak{A}}(\overline{a}) \mid (\overline{a}^{(1)}, \ldots, \overline{a}^{(m)}) \in \alpha\}.$$

At the atomic level, $i = 0$, this representation is sound because the atomic $\tau^{[k]}$-type of $(\overline{a}^{(1)}, \ldots, \overline{a}^{(m)})$ is directly determined by the atomic τ-type of $\overline{a}$. It remains to consider the refinement step — soundness of the representation and PTIME computability in terms of $I_{C^{mk}}(\mathfrak{A})$. Let α be a $\approx_i$-class, $\underline{\alpha}$ its representation. The refinement step is governed by the counting functions

$$\begin{aligned} \nu^{\alpha}_j(\overline{a}) &= \left|\{\overline{b}^{(j)} \mid (\overline{a}^{(1)}, \ldots, \overline{a}^{(m)})\tfrac{\overline{b}^{(j)}}{j} \in \alpha\}\right| \\ &= \left|\{\overline{b}^{(j)} \mid \mathrm{tp}^{C^{mk}}_{\mathfrak{A}}((\overline{a}^{(1)}, \ldots, \overline{a}^{(m)})\tfrac{\overline{b}^{(j)}}{j}) \in \underline{\alpha}\}\right|. \end{aligned}$$

These values clearly only depend on the C^{mk}-type of $\overline{a}$, and they are PTIME computable from $(I_{C^{mk}}(\mathfrak{A}), \underline{\alpha})$. In fact $\nu^{\alpha}_j(\overline{a})$ is the cardinality of a definable k-ary predicate, definable in terms of a union of C^{mk}-equivalence classes that is represented over $I_{C^{mk}}(\mathfrak{A})$ through $\underline{\alpha}$. □

Proof of Theorem 6.16. Let H^3 be a PTIME canonization functor for C^3 on $\mathrm{fin}[\tau^{[k]}]$. By definition H^3 satisfies for all $\mathfrak{A}^{[k]}$: $H^3(\mathfrak{A}^{[k]}) \equiv^{C^3} \mathfrak{A}^{[k]}$. With Lemma 6.22 we conclude that $H^3(\mathfrak{A}^{[k]}) \in \Gamma_k(\mathrm{fin}[\tau])$, since $\mathfrak{A}^{[k]} \in \Gamma_k(\mathrm{fin}[\tau])$ and $\Gamma_k(\mathrm{fin}[\tau])$ is closed under C^3-equivalence. Therefore $H^3(\mathfrak{A}^{[k]}) \simeq \mathfrak{B}^{[k]}$ for some $\mathfrak{B} \in \mathrm{stan}[\tau]$ that is PTIME computable from $H^3(\mathfrak{A}^{[k]})$ by Lemma 6.22. We infer from Proposition 6.21 that $\mathfrak{B} \equiv^{C^k} \mathfrak{A}$. By Lemma 6.23, $H^3(\mathfrak{A}^{[k]})$ and therefore the resulting $\mathfrak{B}$ are fully determined by the C^{3k}-theory of $\mathfrak{A}$. The composite mapping $\Gamma_k^{-1} \circ H^3 \circ \Gamma_k$ is thus seen to satisfy the requirements on $K^{3k,3}$ in the theorem. The statement concerning the L^k rather than the C^k is obtained in exactly the same manner. □

In complete analogy we also prove Theorem 6.17.

Proof of Theorem 6.17. Consider first the case with counting quantifiers. Assume that F^3 is a PTIME inverse for I_{C^3}. For all $\mathfrak{B}$: $F^3(I_{C^3}(\mathfrak{B})) \equiv^{C^3} \mathfrak{B}$. Since membership in $\Gamma_k(\mathrm{fin}[\tau])$ is a C^3-property by Lemma 6.22, F^3 restricts to $\{I_{C^3}(\mathfrak{B}) \mid \mathfrak{B} \in \Gamma_k(\mathrm{fin}[\tau])\}$ such that

$$F^3 : \{I_{C^3}(\mathfrak{B}) \mid \mathfrak{B} \in \Gamma_k(\mathrm{fin}[\tau])\} \longrightarrow \mathrm{stan}[\tau^{[k]}] \cap \Gamma_k(\mathrm{fin}[\tau]).$$

Let I be the mapping I: $I_{C^{3k}}(\mathfrak{A}) \mapsto I_{C^3}(\mathfrak{A}^{[k]})$, which is in PTIME according to Lemma 6.23. The composite mapping $G^{3k,k} = \Gamma_k^{-1} \circ F^3 \circ I$ satisfies the requirement of Theorem 6.17:

$$\forall \mathfrak{A} \quad G^{3k,k}\Big(I_{C^{3k}}(\mathfrak{A})\Big) \equiv^{C^k} \mathfrak{A}.$$

This is because

$$F^3 \circ I(I_{C^{3k}}(\mathfrak{A})) = F^3(I_{C^3}(\mathfrak{A}^{[k]})) \equiv^{C^3} \mathfrak{A}^{[k]} = \Gamma_k(\mathfrak{A}).$$

$\Gamma_k^{-1} \circ F^3 \circ I(I_{C^{3k}}(\mathfrak{A})) \equiv^{C^k} \mathfrak{A}$ now follows with Proposition 6.21.

For the case of L^k one merely checks in addition that the modified notion of PTIME inversion adapted to the I_{L^k} carries over from the corresponding given F^3 to the composite mapping $G^{3k,k}$. Note that the minimal size of structures $\mathfrak{B}$ that are L^3-equivalent with $\mathfrak{A}^{[k]}$ is bounded from above by the k-th power of the size of any structure $\mathfrak{B}'$ that is L^{3k}-equivalent with $\mathfrak{A}$. □

Proof of Lemma 6.22. Obviously any $\mathfrak{A}^{[k]}$ satisfies the following axioms that are all in $L^3_{\omega\omega}$.

(1) $\displaystyle\bigwedge_{i,j} \forall x \forall y \Big(x \stackrel{i,j}{=} y \longleftrightarrow y \stackrel{j,i}{=} x\Big).$

(2) $\displaystyle\forall x \forall y \Big(\bigwedge_{i} x \stackrel{i,i}{=} y \longleftrightarrow x = y\Big).$

(3) $\displaystyle\bigwedge_{i,j,l} \forall x \forall y \forall z \Big(x \stackrel{i,j}{=} y \wedge y \stackrel{j,l}{=} z \longrightarrow x \stackrel{i,l}{=} z\Big).$

$$(4) \quad \bigwedge_{s\subseteq\{1,\dots,k\}} \forall x\forall y\exists z\Big(\bigwedge_{i\in s} z \overset{i,i}{=} x \wedge \bigwedge_{i\notin s} z \overset{i,i}{=} y\Big).$$

Axioms (1) – (4) exclusively concern the equality structure. We add a finite schema of axioms that formalize compatibility conditions between the atomic types as encoded in the P_θ and the equality structure described by the $\overset{i,j}{=}$. Let Ψ be the set of all quantifier free $\tau^{[k]}$-formulae in two variables x and y that are valid in $\Gamma_k(\text{fin}[\tau])$; we can restrict these to some syntactic normal form to keep the set finite without changing its semantics. In fact, the quantifier free kernels of (1) and (2) above are also represented in Ψ. Thus (1) and (2) become redundant when we now further put

$$(5) \quad \forall x\forall y \bigwedge_{\psi\in\Psi} \psi.$$

We first show that

(A) any $(\overset{i,j}{=})$-structure satisfying (1) – (4) is isomorphic with a structure $\big(\{0,\dots,n-1\}^k, (\overset{i,j}{=})\big)$ with the natural interpretation for the $\overset{i,j}{=}$.

The isomorphism is unique up to a permutation of $n = \{0,\dots,n-1\}$.

To prove (A) let $\mathfrak{B} = \big(B, (\overset{i,j}{=}{}^{\mathfrak{B}})\big)$ be a model of (1) – (4). Observe that (1) – (3) imply that the $\overset{i,i}{=}$ are equivalence relations on B whose common refinement is equality. Denote by $[b]_i$ the equivalence class of b with respect to $\overset{i,i}{=}$. It follows from (2) that

$$\begin{aligned} \pi: B &\longrightarrow \textstyle\prod_i B/\overset{i,i}{=} \\ b &\longmapsto ([b]_1,\dots,[b]_k) \end{aligned}$$

is an injection. We show that (4) implies π is surjective. Assume to the contrary that some $([b_1]_1,\dots,[b_k]_k)$ is not in the image of π. Then at least one of

$$\begin{aligned} &\big([b_1]_1,\dots,[b_{k-2}]_{k-2},[b_{k-1}]_{k-1},[b_{k-1}]_k\big) \\ \text{or} \quad &\big([b_1]_1,\dots,[b_{k-2}]_{k-2},[b_k]_{k-1},[b_k]_k\big) \end{aligned}$$

is not in the image of π. Otherwise, choosing pre-images under π of these for x and y and applying (4) with $s = \{1,\dots,k-1\}$, one would get a pre-image of $([b_1]_1,\dots,[b_k]_k)$. Proceeding inductively we obtain that for some b, $([b]_1,\dots,[b]_k)$ is not in the image of π, which is clearly absurd. Therefore π is a bijection. By definition it maps $\overset{i,i}{=}$ to equality in the i-th component.

Finally, the $\overset{i,j}{=}{}^{\mathfrak{B}}$ induce bijections between the different factors $B/\overset{i,i}{=}$. This follows from (3): (3) implies that $\overset{i,j}{=}$ is closed under $\overset{i,i}{=}$ on the left and under $\overset{j,j}{=}$ on the right, so that it factorizes to yield a binary relation between $B/\overset{i,i}{=}$ and $B/\overset{j,j}{=}$. We claim that in this sense it becomes the graph of a bijection. For reasons of symmetry (1) it suffices to show injectivity, or that

$$b_1 \stackrel{i,j}{=} b \wedge b_2 \stackrel{i,j}{=} b' \wedge b \stackrel{j,j}{=} b' \implies b_1 \stackrel{i,i}{=} b_2,$$

which is immediate from (3) and (1). Let $\rho^{ij}\colon B/\stackrel{i,i}{=} \longrightarrow B/\stackrel{j,j}{=}$ be this bijection. Then

$$\begin{aligned} \widehat{\pi}\colon B &\longrightarrow \underbrace{(B/\stackrel{1,1}{=}) \times \cdots \times (B/\stackrel{1,1}{=})}_{k} \\ b &\longmapsto \big(\rho^{i1}([b]_i)\big)_{i=1,\ldots,k} \end{aligned}$$

is an isomorphism between $\mathfrak{B}$ and $(B/\stackrel{1,1}{=})^k$ with the standard interpretation for the $\stackrel{i,j}{=}$. We have shown that (A) holds. The proof also shows that the desired isomorphism with some $\big(\{0,\ldots,n-1\}^k, (\stackrel{i,j}{=})\big)$ is unique up to the choice of an identification of $B/\stackrel{1,1}{=}$ with the appropriate set n. If $\mathfrak{B}$ is itself presented as a standard structure over some $\{0,\ldots,n^k-1\}$, then the natural order on B induces an ordering of $B/\stackrel{1,1}{=}$ which can be used to determine a unique isomorphism of $\mathfrak{B}$ with $\big(\{0,\ldots,n-1\}^k, (\stackrel{i,j}{=})\big)$. This isomorphism is constructible in PTIME.

Assume now that $\mathfrak{B}$ carries interpretations for the P_θ and is a model also of (5). For the full claim of the lemma it remains to translate the information in the P_θ to a τ-interpretation over n.

Let now $\pi\colon \big(\mathfrak{B} \restriction (\stackrel{i,j}{=})\big) \longmapsto \big(\{0,\ldots,n-1\}^k, (\stackrel{i,j}{=})\big)$ be an isomorphism with the standard model of the equality part. (5) implies in particular that

(a) the $P_\theta^{\mathfrak{B}}$ form a partition of B. Introduce the mapping $\Theta\colon B \to \mathrm{Atp}(\tau;k)$ which sends b to that θ with $b \in P_\theta^{\mathfrak{B}}$.
(b) $\pi(b) = (m_1,\ldots,m_k)$ implies that the equality type of $(m_1,\ldots,m_k)$ is as prescribed in $\Theta(b)$.
(c) if $\pi(b) = (m_1,\ldots,m_k)$ and $\pi(b') = (m_1',\ldots,m_k')$, then the instantiations $\big(\Theta(b)\big)[m_1,\ldots,m_k]$ and $\big(\Theta(b')\big)[m_1',\ldots,m_k']$ are consistent, i.e. respect the equality type of the tuple $(m_1,\ldots,m_k,m_1',\ldots,m_k')$.

It follows that $\{0,\ldots,n-1\}$ can be expanded to a τ-structure $\mathfrak{A}$ in a unique and consistent way by the stipulation that $\mathrm{atp}_{\mathfrak{A}}\big(\pi(b)\big) = \Theta(b)$. Thus π becomes an isomorphism between $\mathfrak{B}$ and $\mathfrak{A}^{[k]}$. $\mathfrak{A}$ can obviously be computed from π and $\mathfrak{B}$ in polynomial time, so that the second claim of Lemma 6.22 also follows. □

It is interesting to note that in the above $L^3_{\omega\omega}$-axiomatization three variables are necessary for the transitivity conditions (3) and the sentence (4), which ensures surjectivity of π. A condition to the effect of (4) can actually also be formalized in $C^2_{\infty\omega}$. Let χ_n be the sentence

$$\exists^{=m} x\, x = x \;\wedge\; \bigwedge_i \forall x\, \exists^{=s} y\; x \stackrel{i,i}{=} y$$

with $m = n^k$, $s = n^{k-1}$. Then (4) above can be replaced by $\bigvee_n \chi_n$. One obtains an axiomatization of $\Gamma_k\big(\mathrm{fin}[\tau]\big)$ in $C^3_{\infty\omega}$ that uses three variables only in the transitivity conditions for the $\stackrel{i,j}{=}$ in (3).

6.4.2 Remarks on Further Reduction

This aside is of a more technical nature. The proofs of Theorems 6.16 and 6.17 given in Section 6.4.1 need seemingly weaker assumptions than full PTIME canonization (or inversion of the invariants) in the three variable cases. We strengthen the formulation accordingly in this section. As all these considerations remain hypothetical — we have no well founded conjecture whether PTIME $\cap\, C^{\omega}_{\infty\omega}$ or PTIME $\cap\, L^{\omega}_{\infty\omega}$ can indeed be captured — the interest in these ramifications mainly is a technical one. We explicitly address the C^k in this aside, but once more everything translates to the L^k. Consider the situation with respect to canonization. The proof of Theorem 6.16 rests on the existence of PTIME computable functors

$$H_0\colon \mathrm{fin}[\tau^{[k]}] \longrightarrow \mathrm{stan}[\tau^{[k]}] \cap \Gamma_k(\mathrm{fin}[\tau]) \tag{6.1}$$

such that for all $\mathfrak{C}$ and $\mathfrak{C}'$ in the domain of H_0:

$$\begin{aligned} &H_0(\mathfrak{C}) \equiv^{C^2} \mathfrak{C} \\ &\mathfrak{C} \equiv^{C^m} \mathfrak{C}' \rightarrow H_0(\mathfrak{C}) = H_0(\mathfrak{C}') \end{aligned} \tag{6.2}$$

for some m. From these we obtain 'weak canonization functors' $K^{mk,k}$ from $\mathrm{fin}[\tau]$ to $\mathrm{stan}[\tau]$ that satisfy

$$\begin{aligned} &\forall\mathfrak{A} \quad K^{mk,k}(\mathfrak{A}) \equiv^{C^k} \mathfrak{A}, \\ &\forall\mathfrak{A}\forall\mathfrak{A}' \quad \mathfrak{A} \equiv^{C^{mk}} \mathfrak{A}' \rightarrow K^{mk,k}(\mathfrak{A}) = K^{mk,k}(\mathfrak{A}'). \end{aligned}$$

In the proof of Theorem 6.16 we have explicitly used this construction for a proper canonization functor H for $C^3_{\infty\omega}$ in place of H_0 and with $m = 3$. Note that in this special case both $H(\mathfrak{A}^{[k]}) \in \Gamma_k(\mathrm{fin}[\tau])$ and $H(\mathfrak{A}^{[k]}) \equiv^{C^2} \mathfrak{A}^{[k]}$ are consequences of the stronger requirement that

$$H(\mathfrak{A}^{[k]}) \equiv^{C^3} \mathfrak{A}^{[k]}.$$

Surprisingly, a twofold application of the reduction schema leads to a further reduction in the assumptions expressed in equations 6.1 and 6.2. One need only assume the existence of such H_0 for $k = 3$. In particular this amounts to a reduction to vocabularies $\tau^{[3]}$ with a fixed set of binary relations $(\overset{i,j}{=})_{1\leqslant i,j\leqslant 3}$ and only unary predicates besides.

Proposition 6.24. *Assume that for each τ there is a* PTIME *functor*

$$H\colon \mathrm{fin}\big[(\tau^{[k]})^{[3]}\big] \longrightarrow \mathrm{stan}\big[(\tau^{[k]})^{[3]}\big] \cap \Gamma_3\big(\mathrm{fin}\big[\tau^{[k]}\big]\big)$$

such that for all $\mathfrak{C}, \mathfrak{C}' \in \mathrm{fin}[(\tau^{[k]})^{[3]}]$ and some fixed m:

$$\begin{aligned} &H(\mathfrak{C}) \equiv^{C^2} \mathfrak{C} \\ &\mathfrak{C} \equiv^{C^m} \mathfrak{C}' \rightarrow H(\mathfrak{C}) = H(\mathfrak{C}'). \end{aligned}$$

Then $H_0 := \Gamma_3^{-1} \circ H \circ \Gamma_3$ *satisfies the conditions in equations 6.1 and 6.2, with* $3m$ *in place of* m *and consequently* $K^{3mk,k} := \Gamma_k^{-1} \circ \Gamma_3^{-1} \circ H \circ \Gamma_3 \circ \Gamma_k$ *provides 'weak canonization' for* $C^k_{\infty\omega}$ *in the sense of Theorem 6.16 on* fin$[\tau]$*:*

$$\begin{aligned} &\forall \mathfrak{A} \quad K^{3mk,k}(\mathfrak{A}) \equiv^{C^k} \mathfrak{A}, \\ &\forall \mathfrak{A} \forall \mathfrak{A}' \quad \mathfrak{A} \equiv^{C^{3mk}} \mathfrak{A}' \rightarrow K^{3mk,k}(\mathfrak{A}) = K^{3mk,k}(\mathfrak{A}'). \end{aligned}$$

Sketch of Proof. The crucial observation is that $H_0 := \Gamma_3^{-1} \circ H \circ \Gamma_3$ is well defined, has image in $\Gamma_k\big(\text{fin}[\tau]\big)$ and satisfies for all $\mathfrak{A} \in \text{fin}[\tau]$: $H_0\big(\mathfrak{A}^{[k]}\big) \equiv^{C^2} \mathfrak{A}^{[k]}$.

Let $\mathfrak{A} \in \text{fin}[\tau]$, $\mathfrak{A}^{[k]}$ its k-th power. Let $H\big(\big(\mathfrak{A}^{[k]}\big)^{[3]}\big) \simeq \mathfrak{B}^{[3]}$. By Proposition 6.21 we know that $\big(\mathfrak{A}^{[k]}\big)^{[3]} \equiv^{C^2} \mathfrak{B}^{[3]}$ implies that $\mathfrak{A}^{[k]} \equiv^{C^3} \mathfrak{B}$. By Lemma 6.22, $\mathfrak{B}$ therefore is itself a k-th power. A further application of Proposition 6.21 yields $\Gamma_k^{-1}(\mathfrak{B}) \equiv^{C^k} \mathfrak{A}$. As $H(\mathfrak{C})$ only depends on the C^m-theory of $\mathfrak{C}$ by assumption, it follows that $\big[H \circ \Gamma_3\big](\mathfrak{A})$ is determined by the C^{3m}-theory of $\mathfrak{A}$ and finally that $\big[H \circ \Gamma_3 \circ \Gamma_k\big](\mathfrak{A})$ is fully determined by the C^{3mk}-theory of $\mathfrak{A}$ (compare Lemma 6.23). □