7. Canonization for Two Variables

In this chapter we prove that both $L^2_{\infty\omega}$ and $C^2_{\infty\omega}$ admit PTIME canonization. We do so by exhibiting PTIME inverses for I_{L^2} and I_{C^2}. The inversion for I_{L^2} is even PTIME in terms of the size of the I_{L^2}, a phenomenon that we know to be peculiar to the two variable case. These are the main theorems:

Theorem 7.1. *I_{L^2} admits* PTIME *inversion in the strong sense that for each finite relational τ there is a* PTIME *functor $F\colon \{I_{L^2}(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{fin}[\tau]\} \to \mathrm{stan}[\tau]$, which is an inverse for I_{L^2}:*

$$\forall\mathfrak{A} \quad F\big(I_{L^2}\mathfrak{A}\big) \equiv^{L^2} \mathfrak{A}.$$

It follows that

(i) the range of I_{L^2} can be recognized in PTIME.
(ii) $L^2_{\infty\omega}$ admits PTIME *canonization.*
(iii) PTIME $\cap\, L^2_{\infty\omega}$ *is recursively enumerable (has a recursive presentation).*
(iv) PTIME $\cap\, L^2_{\infty\omega} \equiv \mathrm{FP}\big(I_{L^2}\big) \equiv \mathrm{PTIME}\big(I_{L^2}\big)$.

Compare the general Theorems 6.11 and 6.14 for (ii) and (iii). (i) is obvious: for $\mathfrak{I}$ of the format of an L^2-invariant, $\mathfrak{I} \in \{I_{L^2}(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{fin}[\tau]\}$ if and only if $F(\mathfrak{I}) \in \mathrm{fin}[\tau]$ and $I_{L^2}\big(F(\mathfrak{I})\big) = \mathfrak{I}$. (i) and the strong form of (iv) (if compared to the statement of Theorem 6.14) are consequences of polynomiality of F in the usual sense.

Theorem 7.2. *I_{C^2} admits* PTIME *inversion. For each finite relational τ there is a* PTIME *functor $F\colon \{I_{C^2}(\mathfrak{A}) \mid \mathfrak{A} \in \mathrm{fin}[\tau]\} \to \mathrm{stan}[\tau]$, which is an inverse for I_{C^2}:*

$$\forall\mathfrak{A} \quad F\big(I_{C^2}\mathfrak{A}\big) \equiv^{C^2} \mathfrak{A}.$$

It follows that

(i) the range of I_{C^2} can be recognized in PTIME.
(ii) $C^2_{\infty\omega}$ admits PTIME *canonization.*
(iii) PTIME $\cap\, C^2_{\infty\omega}$ *is recursively enumerable (has a recursive presentation).*
(iv) PTIME $\cap\, C^2_{\infty\omega} \equiv \mathrm{FP}\big(I_{C^2}\big) \equiv \mathrm{PTIME}\big(I_{C^2}\big)$.

The construction of the inverses is reduced to a combinatorial problem that only deals with the abstract information about the corresponding two-pebble games as represented in the invariants. The relational information encoded in the invariants through the identification of atomic types is at first suppressed for this purpose. The reduced invariants, stripped of the relational particulars, are what we shall call *game tableaux.*

- These game tableaux are introduced in Section 7.1. The inversion problems for I_{L^2} and I_{C^2} are reduced to the problem of constructing realizations for tableaux.
- In Section 7.2 such realizations are constructed in the case of C^2.
- Section 7.3 deals with the corresponding constructions in the case of L^2.

7.1 Game Tableaux and the Inversion Problem

For the start $L^2_{\infty\omega}$ and $C^2_{\infty\omega}$ can be treated in parallel. Recall from Definitions 3.3 and 3.11 the format of the invariants I_{C^2} and I_{L^2}. The special situation in dimension two allows for certain simplifications in their presentation. Consider I_{C^2} first. In the original format:

$$I_{C^2}(\mathfrak{A}) \;=\; \Big(A^2/{\equiv^{C^2}}, \leqslant, (P_\theta)_{\theta\in\mathrm{Atp}(\tau;2)}, (E_j)_{j=1,2}, (S_\rho)_{\rho\in S_2}; (\nu_j)_{j=1,2}\Big).$$

As pointed out in the general case, it suffices to retain one of the E_j and ν_j each, since the other one remains definable with the help of $S_{(1,2)}$, the encoding of the permutation that exchanges first and second component. This permutation is the only member of S_2 apart from the identity. In the following we denote by T (for transposition) both this exchange of first and second component as a member of S_2 and its operation on the elements of I_{C^2}. The graph of this operation, $S_{(1,2)}$, is also denoted T. This is not likely to cause any confusion, since the transition between these representations is trivial. Retaining $E := E_2$ and $\nu := \nu_2$, we get:

$$\begin{array}{rcrcl} E_1 & = & E^T & := & \big\{(\alpha,\alpha') \mid (T\alpha, T\alpha') \in E\big\}, \\ \nu_1 & = & \nu^T & := & \nu \circ T. \end{array}$$

We separate the equality type information from the remaining relational atomic information in the P_θ by putting

$$\Delta := \big\{\alpha \mid (a_1,a_2) \in \alpha \Rightarrow a_1 = a_2\big\}.$$

For notational convenience finally, the partition of the universe into the P_θ is replaced by a function $\Theta\colon A^2/\equiv^{C^2} \to \mathrm{Atp}(\tau;2)$. We thus obtain the following format for the I_{C^2}, which is obviously interdefinable at first-order level with the former one:

$$I_{C^2}(\mathfrak{A}) \;=\; \Big(A^2/\equiv^{C^2}, \leqslant, E, T, \Delta; \Theta, \nu\Big). \tag{7.1}$$

The same modifications apply to I_{L^2}:

$$I_{L^2}(\mathfrak{A}) \;=\; \Big(A^2/\equiv^{L^2}, \leqslant, E, T, \Delta; \Theta\Big). \tag{7.2}$$

Proviso. For the purposes of this chapter we fix the special format for the two-variable invariants according to equations 7.1 and 7.2 above. We regard both I_{C^2} and I_{L^2} as (standard representations of) ordered weighted κ-structures, where $\kappa := \{\leqslant, E, T, \Delta\}$.

We collect a few obvious facts about the κ-reducts of two-variable invariants, no matter whether I_{C^2} or I_{L^2}, in the following lemma.

Lemma 7.3. *Let $\mathfrak{Q} = (Q, \leqslant, E, T, \Delta)$ be the κ-reduct of some $I_{C^2}(\mathfrak{A})$ or $I_{L^2}(\mathfrak{A})$. Then $\mathfrak{Q}$ satisfies the following:*

(i) E is an equivalence relation on Q.
(ii) T is (the graph of) an involutive function from Q to Q: $T \circ T = \mathrm{id}_Q$.
(iii) Δ consists of points fixed under T: $T \restriction \Delta = \mathrm{id}_\Delta$.
(iv) each E-class contains exactly one element from Δ.

Proof. (i) – (iii) are obvious on the basis of the definitions. Note in connection with (i) that for the underlying invariant $I_{C^2}(\mathfrak{A})$ or $I_{L^2}(\mathfrak{A})$ an E-class exactly corresponds to the type of a single element of $\mathfrak{A}$. This may be seen as follows. Let $\mathcal{L} = C^2_{\infty\omega}$ or $L^2_{\infty\omega}$, $I_\mathcal{L}$ the corresponding invariant, $\mathfrak{Q} = I_\mathcal{L}(\mathfrak{A}) \restriction \kappa$. Let q be the E-class of an element $\alpha \in Q$. Fix some (a_1, a_2) such that $\alpha = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_2)$. Then by definition q consists of exactly those α' with $\alpha' = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_2')$ for some $a_2' \in A$. Let $\beta = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1)$. β is fully determined by q since it exactly consists of all those formulae $\varphi(x_1) \in \mathcal{L}[\tau]$ that are members of all $\alpha' \in q$. Conversely, q itself is completely determined by β, since $\alpha' \in Q$ if and only if $\exists x_2 \varphi_{\alpha'}(x_1, x_2) \in \beta$, for some formula $\varphi_{\alpha'}(x_1, x_2)$ that isolates α'.

For (iv) first observe that there must be an element from Δ in each E-class of a real invariant. If $\alpha = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_2)$, then $\mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_1)$ is in Δ and E-related with α. For uniqueness as claimed in (iv) consider $\delta_1, \delta_2 \in \Delta$ and assume that δ_1 and δ_2 are E-related. $\delta_1 = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_1)$ for some a_1, and by E-relatedness there must be some a_2 such that $\delta_2 = \mathrm{tp}^\mathcal{L}_\mathfrak{A}(a_1, a_2)$. Since $\delta_2 \in \Delta$, $a_1 = a_2$ and therefore $\delta_1 = \delta_2$. □

The following definition introduces the term *game tableaux* for those κ-structures that are candidates for the relational parts of two-variable invariants according to the last lemma. Note that Θ, the assignment of relational atomic types, is not made part of the game tableaux.

Definition 7.4. *A finite κ-structure $\mathfrak{Q} = (Q, \leqslant, E, T, \Delta)$ is called a* game tableau *if and only if $\leqslant$ is a linear ordering on Q and $\mathfrak{Q}$ satisfies conditions (i) – (iv) of Lemma 7.3. A* weighted game tableau *is a game tableau $\mathfrak{Q}$ together with a weight function $\nu: Q \to \omega \setminus \{0\}$.*

The size of a game tableau $\mathfrak{Q}$ is its size as a relational structure, i.e. the size of its universe Q. The size of a weighted game tableau $(\mathfrak{Q}; \nu)$ is taken to be $\sum_{\alpha \in Q} \nu(\alpha)$. These conventions ensure that the size of the original invariant is polynomially related to the size of the abstracted tableau.

As the $\mathfrak{Q}$ and $(\mathfrak{Q}; \nu)$ are linearly ordered, we may think of them as standard objects. The standardization is then implicitly assumed to be the same as for the I_{L^2} and I_{C^2} in their new format.

Note that the class of game tableaux is first-order definable. Game tableaux and weighted game tableaux are recognizable in LOGSPACE. The following lemma isolates some obvious conditions on the function Θ that in real invariants associates relational types with the elements of the invariant. The proof is immediate and similar to that of Lemma 7.3 above.

Lemma 7.5. *Let $\mathfrak{Q}$ be the κ-reduct of $I_{C^2}(\mathfrak{A})$ or $I_{L^2}(\mathfrak{A})$ for some $\mathfrak{A}$, $\Theta: Q \rightarrow$* $\mathrm{Atp}(\tau; 2)$ *the mapping that associates the relational atomic types with the elements of the invariant. Then Θ satisfies the following conditions:*

(i) if $\alpha \in \Delta$ then $\Theta(\alpha)$ is the type of an identity pair: $x_1 = x_2 \in \Theta(\alpha)$.
(ii) for all $\alpha \in Q$, $\Theta\big(T(\alpha)\big)$ is the atomic type obtained from $\Theta\big(\alpha\big)$ by exchanging x_1 and x_2 in all formulae.
(iii) if δ is the unique element of Δ that is in the E-class of α, then $\Theta(\delta)$ contains all formulae $\varphi(x_1)$ from $\Theta(\alpha)$.

Note that the syntactic conditions in (i), (ii) and (iii) completely determine $\Theta\big(T(\alpha)\big)$ in terms of $\Theta(\alpha)$ in (ii) and $\Theta(\delta)$ in terms of $\Theta(\alpha)$ in (iii).

Definition 7.6. *Let $\mathfrak{Q}$ be a game tableau, Θ a function from its domain to* $\mathrm{Atp}(\tau; 2)$ *for some τ. Θ is a* good extension *of $\mathfrak{Q}$ if conditions (i) – (iii) of Lemma 7.5 are satisfied.*

It can be checked in LOGSPACE whether Θ is a good extension of $\mathfrak{Q}$.

The inversion of an invariant asks for the construction of a relational structure over some n such that the types of pairs in this structure fit the specifications laid down in the given invariant. We first approach this problem at the level of the underlying game tableaux or weighted game tableaux — the relational atomic types, as encoded in the Θ, are disregarded at first. Correspondingly, the result of this approach is somewhat less than a relational structure. We shall call it a *realization* of the given game tableau. It turns out that these realizations govern the combinatorial pattern of relational structures to such an extent that the plain relational information in Θ need only be added in later. Formally we describe the desired realizations as mappings that associate pairs over some standard domain with elements of the game tableau. The intention is that — once we also plug in relational information — this mapping will actually be the projection sending pairs to their types.

Definition 7.7. *Let $\mathfrak{Q}$ be a game tableau. A surjective mapping $\pi: n \times n \to Q$ is called a* realization of $\mathfrak{Q}$ over the standard domain n, *if the following conditions are satisfied, where m_1, m_2 range over the elements of n:*

(i) π respects the diagonal: $\pi(m_1, m_2) \in \Delta$ if and only if $m_1 = m_2$.
(ii) π respects T: $\pi(m_1, m_2) = T\big(\pi(m_2, m_1)\big)$, i.e. π commutes with T.
(iii) π respects E: The E-class of $\pi(m_1, m_2)$ is the set of all $\pi(m_1, m_2')$ for $m_2' \in n$.

If $\nu: Q \to \omega \setminus \{0\}$ is a weight function on $\mathfrak{Q}$, then we further say that π realizes the weighted tableau $(\mathfrak{Q}; \nu)$ if also

(iv) π is compatible with ν:
$\nu\big(\pi(m_1, m_2)\big) = \big|\{m_2' \in n \mid \pi(m_1, m_2') = \pi(m_1, m_2)\}\big|$.

Obviously the definition states a number of conditions that are always satisfied in case that $\mathfrak{Q}$ (and ν) are derived from a real invariant of a structure over n and if π is the natural projection sending pairs of elements to their types. We state this fact as a lemma; the proof is trivial.

Lemma 7.8. *Let $\mathfrak{A} \in \mathrm{stan}[\tau]$ be a τ-structure over universe n. Let $\mathcal{L} = C^2_{\infty\omega}$ or $L^2_{\infty\omega}$, $I_{\mathcal{L}}(\mathfrak{A})$ the corresponding invariant. Let $\mathfrak{Q}$ be the induced game tableau, so that $Q = \mathrm{Tp}^{\mathcal{L}}(\mathfrak{A}; 2)$. Put*

$$\begin{array}{rccl} \pi: & n \times n & \longrightarrow & Q \\ & (m_1, m_2) & \longmapsto & \mathrm{tp}^{\mathcal{L}}_{\mathfrak{A}}(m_1, m_2). \end{array}$$

Then π is a realization of the tableau $\mathfrak{Q}$. In case $\mathcal{L} = C^2_{\infty\omega}$ and if ν is the weight function of $I_{C^2}(\mathfrak{A})$, π is a realization of the weighted tableau $(\mathfrak{Q}; \nu)$.

A realization of a game tableau over n, together with attributions of atomic τ-types (in the form a some good extension) uniquely determines a τ-structure with domain n. Let $\mathfrak{Q}$ be realized by $\pi: n \times n \to Q$ and let $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ be good in the sense of Definition 7.6. Assume first that τ contains no relation symbols of arity greater than 2. Then there is a unique structure $\mathfrak{A}(\pi, \Theta) \in \mathrm{stan}[\tau]$ over n for which

$$\forall m_1 \forall m_2 \;\; \mathrm{atp}_{\mathfrak{A}}(m_1, m_2) = \Theta\big(\pi(m_1, m_2)\big).$$

Uniqueness is obvious. For the existence claim one has to check that the conditions expressed in the above equations are compatible. The requirements for realizations and for good Θ are designed just to guarantee this compatibility. For instance if $m_1 = m_2$, we have, by a corresponding condition on realizations, that $\pi(m_1, m_1) \in \Delta$, whence it follows that $\Theta\big(\pi(m_1, m_1)\big)$ is an atomic type of an identity pair. For any $m_1, m_2 \in n$, compatibility of $\Theta\big(\pi(m_1, m_2)\big)$ with $\Theta\big(\pi(m_1, m_1)\big)$ follows from the fact that $\pi(m_1, m_1)$ must be the unique element of Δ in the E-class of $\pi(m_1, m_2)$, since π respects E. But then $\Theta\big(\pi(m_1, m_1)\big)$ corresponds to the restriction of $\Theta\big(\pi(m_1, m_2)\big)$ to the first component as Θ is good.

In order to extend the definition of $\mathfrak{A}(\pi, \Theta)$ in a well-defined way to the general case in which relation symbols of arity greater than 2 are admitted in τ, we stipulate that no tuple involving more than two distinct components is put into the interpretation of such relations.

Definition 7.9. *For a realization $\pi: n \times n \to Q$ of a game tableau $\mathfrak{Q}$ and good $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ let $\mathfrak{A}(\pi, \Theta)$ be the unique τ-structure over standard universe n induced by π and Θ as described above.*

Note that $\mathfrak{A}(\pi, \Theta)$ is constructible from π and Θ in PTIME.

The following is an obvious statement to the effect that in the intended case — the case that all the data are obtained from a real structure over some standard universe — $\mathfrak{A}(\pi, \Theta)$ essentially reproduces that original structure.

Lemma 7.10. *Let $\mathfrak{A} \in \mathrm{stan}[\tau]$, $\mathcal{L} = C^2_{\infty\omega}$ or $L^2_{\infty\omega}$, $\mathfrak{Q}$ the game tableau induced by $I_{C^2}(\mathfrak{A})$ or $I_{L^2}(\mathfrak{A})$, respectively. Let $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ be the good extension induced by the invariant itself. Let $\pi: n \times n \longrightarrow Q$ be the realization that is the natural projection $\pi: A^2 \to A^2/\equiv^{\mathcal{L}}$. If τ contains no relation symbols of arity greater than 2 then $\mathfrak{A}(\pi, \Theta) = \mathfrak{A}$. Otherwise $\mathfrak{A}$ and $\mathfrak{A}(\pi, \Theta)$ agree on all atoms involving at most two elements, so that at least $\mathfrak{A} \equiv^{\mathcal{L}} \mathfrak{A}(\pi, \Theta)$.*

The following proposition is crucial for showing that inversion for the I_{C^2} and I_{L^2} reduces to the construction of realizations for (weighted) tableaux.

Proposition 7.11. *Let $\mathfrak{Q}$ be a game tableau, $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ a good extension of $\mathfrak{Q}$.*

(i) If π and π' are any two realizations of $\mathfrak{Q}$ then $\mathfrak{A}(\pi, \Theta) \equiv^{L^2} \mathfrak{A}(\pi', \Theta)$.
(ii) If $\nu: Q \to \omega \setminus \{0\}$ is a weight function on $\mathfrak{Q}$ and π and π' are any two realizations of the weighted tableau $(\mathfrak{Q}; \nu)$ then $\mathfrak{A}(\pi, \Theta) \equiv^{C^2} \mathfrak{A}(\pi', \Theta)$.

Proof. Consider (ii), the case of C^2. It has to be shown that $\mathfrak{A} := \mathfrak{A}(\pi, \Theta)$ and $\mathfrak{A}' := \mathfrak{A}(\pi', \Theta)$ satisfy exactly the same C^2-types. Using the game characterization for C^2-equivalence, Theorem 2.2, we show that player **II** has a strategy to maintain the condition that $\pi(a_1, a_2) = \pi'(a_1', a_2')$ throughout all stages $(\mathfrak{A}, (a_1, a_2); \mathfrak{A}', (a_1', a_2'))$ in the infinite C^k-game. This is the natural condition since realizations are modelled to describe the projections to the C^2-types. This condition is also sufficient for a strategy in the game since the atomic types of pairs over $\mathfrak{A}$ and $\mathfrak{A}'$ are determined by $\Theta \circ \pi$ and $\Theta \circ \pi'$ respectively.

Assume that $\pi(a_1, a_2) = \pi'(a_1', a_2') = \alpha_0$ in the current position. Everything is explicitly symmetric with respect to $\mathfrak{A}$ and $\mathfrak{A}'$ and implicitly also with respect to first or second component, since realizations and good extensions respect T. Let therefore without loss of generality player **I** choose pebble 2 and put forward the challenge $B \subseteq A$. Let $q \subseteq Q$ be the E-class of α_0. For each $\alpha \in q$, let $B_\alpha := \{b \in B \mid \pi(a_1, b) = \alpha\}$. It follows that B is the disjoint union of the B_α — only $\alpha \in q$ need be considered since π respects E. Since π is also compatible with ν, $|B_\alpha| \leqslant \nu(\alpha)$. π' respects E and ν as well so that in $\mathfrak{A}'$ there are disjoint subsets $B'_\alpha \subseteq \{b' \in A' \mid \pi'(a'_1, b') = \alpha\}$ with $|B_\alpha| = |B'_\alpha|$. In fact, for each $\alpha \in q$ we have that the size of $\{b' \in A' \mid \pi'(a'_1, b') = \alpha\}$ must be $\nu(\alpha)$. Let **II** respond with $B' := \bigcup_{\alpha \in q} B'_\alpha$. In the second exchange of this round **I** now chooses $b' \in B'_\alpha$ for some $\alpha \in q$, so that **II** can answer with any $b \in B_\alpha$ and the desired equality $\pi(a_1, b) = \pi'(a'_1, b') = \alpha$ is maintained. □

Together with Lemma 7.10 this proposition yields the main preparatory result for the construction of inverses: full reduction of the inversion problem to that of finding realizations. We give separate statements for $L^2_{\infty\omega}$ and $C^2_{\infty\omega}$.

Theorem 7.12. *Let $\mathfrak{Q}$ be a game tableau, $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ a function. The following are equivalent:*

(i) $(\mathfrak{Q}; \Theta) = I_{L^2}(\mathfrak{A})$ *for some* $\mathfrak{A} \in \mathrm{fin}[\tau]$.
(ii) Θ *is a good extension of* $\mathfrak{Q}$ *and there is a realization* π *of* $\mathfrak{Q}$ *such that* $I_{L^2}(\mathfrak{A}(\pi, \Theta)) = (\mathfrak{Q}; \Theta)$.
(iii) Θ *is a good extension of* $\mathfrak{Q}$*, there is a realization of* $\mathfrak{Q}$*, and for all realizations* π *of* $\mathfrak{Q}$: $I_{L^2}(\mathfrak{A}(\pi, \Theta)) = (\mathfrak{Q}; \Theta)$.

Theorem 7.13. *For a weighted game tableau* $(\mathfrak{Q}; \nu)$ *and a function* $\Theta: Q \to \mathrm{Atp}(\tau; 2)$ *the following are equivalent:*

(i) $(\mathfrak{Q}; \Theta, \nu) = I_{C^2}(\mathfrak{A})$ *for some* $\mathfrak{A} \in \mathrm{fin}[\tau]$.
(ii) Θ *is a good extension of* $\mathfrak{Q}$ *and there is a realization* π *of* $(\mathfrak{Q}; \nu)$ *such that* $I_{C^2}(\mathfrak{A}(\pi, \Theta)) = (\mathfrak{Q}; \Theta, \nu)$.
(iii) Θ *is a good extension of* $\mathfrak{Q}$*, there is a realization of* $(\mathfrak{Q}; \nu)$*, and for all realizations* π *of* $(\mathfrak{Q}; \nu)$: $I_{C^2}(\mathfrak{A}(\pi, \Theta)) = (\mathfrak{Q}; \Theta, \nu)$.

Proof. The proof is indicated for the case of C^2: (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i) is obvious. Assume (i). Without loss of generality $\mathfrak{A} \in \mathrm{stan}[\tau]$. Then the natural projection $\pi: A^2 \to A^2 /{\equiv^{C^2}}$ yields a realization, see Lemma 7.8. Any two realizations of $(\mathfrak{Q}; \nu)$ lead to C^2-equivalent structures $\mathfrak{A}(\pi, \Theta)$ by Proposition 7.11; we therefore get (iii). □

These theorems reduce the proof of the main theorems on PTIME inversion for I_{L^2} and I_{C^2} to the following claims. Recall for complexity considerations that the size of a tableau $\mathfrak{Q}$ is the size of its universe Q as usual, while the size of a weighted tableau $(\mathfrak{Q}; \nu)$ is $\sum_{\alpha \in Q} \nu(\alpha)$.

Theorem 7.14. *There are* PTIME *algorithms* $\mathcal{A}$ *and* $\mathcal{A}^*$ *defined on all* κ*-structures* $\mathfrak{Q}$*, respectively on all* κ*-structures with positive weights* $(\mathfrak{Q}; \nu)$*, such that*

(a) if $\mathfrak{Q}$ *is a game tableau that admits any realization then* $\mathcal{A}$ *applied to* $\mathfrak{Q}$ *yields a realization of* $\mathfrak{Q}$.
(b) if $(\mathfrak{Q}; \nu)$ *is a weighted game tableau that admits any realization then* $\mathcal{A}^*$ *applied to* $(\mathfrak{Q}; \nu)$ *yields a realization of* $(\mathfrak{Q}; \nu)$.

Such algorithms provide the basis for Theorems 7.1 and 7.2. We sketch an algorithm F as required in Theorem 7.1 with respect to $L^2_{\infty\omega}$. The case of $C^2_{\infty\omega}$ is entirely analogous. The input is a structure $\mathfrak{I} = (\mathfrak{Q};\Theta)$ of the format of an $I_{L^2}(\mathfrak{A})$. The following diagram describes the desired algorithm:

- $\mathfrak{I} = (\mathfrak{Q};\Theta)$; check whether $\mathfrak{Q}$ is a game tableau and whether Θ is a good extension of $\mathfrak{Q}$

  > if $\mathfrak{Q}$ is not a game tableau, or if Θ is not a good extension of $\mathfrak{Q}$, then $\mathfrak{I} \notin \mathrm{range}(I_{L^2})$

- In the positive case apply $\mathcal{A}$ to $\mathfrak{Q}$ and check whether the output $\mathcal{A}(\mathfrak{Q})$ is a realization of $\mathfrak{Q}$

  > if the output is not a realization of $\mathfrak{Q}$, then $\mathfrak{I} \notin \mathrm{range}(I_{L^2})$

- Let in the positive case $\pi = \mathcal{A}(\mathfrak{Q})$ be that realization
- Construct $\mathfrak{A}(\pi,\Theta) \in \mathrm{stan}[\tau]$
- Compute $I_{L^2}(\mathfrak{A}(\pi,\Theta))$ and compare with $\mathfrak{I}$

  > if $I_{L^2}(\mathfrak{A}(\pi,\Theta)) \neq \mathfrak{I}$, then $\mathfrak{I} \notin \mathrm{range}(I_{L^2})$

- In the positive case output $F(\mathfrak{I}) := \mathfrak{A}(\pi,\Theta)$.

Correctness essentially depends on Theorem 7.12, which says that any realization of the game tableau leads to a successful construction of an inverse to the invariant if there is any! The rest of this chapter is devoted to the proof of Theorem 7.14.

7.1.1 Modularity of Realizations

This section exhibits an important modularity property of the game tableaux that facilitates the construction of realizations. The overall problem can be decomposed into simpler subproblems, whose solutions form the building blocks for the desired realization.

Definition 7.15. *Let $\mathfrak{Q}$ be a game tableau. We enumerate the E-classes as $q_1,\dots,q_l$. Here $l = |Q/E|$ and the ordering is that induced by $\leqslant^{\mathfrak{Q}}$ in terms of $\leqslant$-least elements of the classes.*

(i) We denote by δ_i the unique element of $q_i \cap \Delta$.
(ii) Let $q_i^T := \{T\alpha \mid \alpha \in q_i\}$.
(iii) Let $q_{ij} := q_i \cap q_j^T$.

Note that the q_i^T are the equivalence classes with respect to E^T, defined by $E^T = \{(\alpha, \alpha') \mid (T\alpha, T\alpha') \in E\}$. For real invariants $E^T = E_1$ is accessibility via a move in the first component. Note also that $T(q_{ij}) = q_{ji}$ and that therefore $q_{ij} \cap q_{ji} = \emptyset$ unless $i = j$; in this case, $\delta_i \in q_{ii} \neq \emptyset$. Simple examples show, however, that by no means need $q_{ii} = \{\delta_i\}$. Consider a directed cycle of length 4 as a graph. For the associated C^2-invariant there is only one type in Δ, E and E^T are both trivial, but there are 4 different C^2-types.

Fig. 7.1

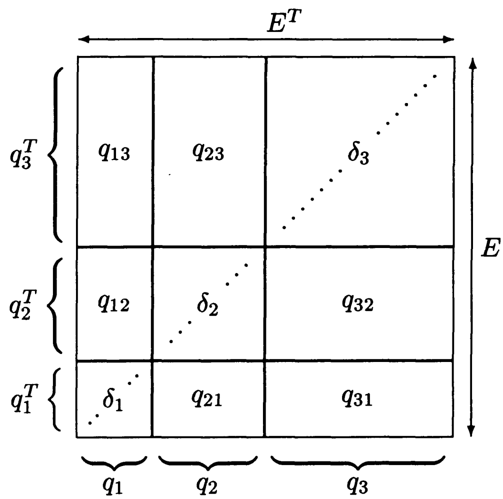

The following characterization of the q_{ij} is technically very useful, for a pictorial presentation see Figure 7.1. In Figure 7.1 the fine structure of Ω is depicted as projected onto some $n \times n$ square that would be a realization.

Lemma 7.16. *Let Ω be a game tableau, π be a realization of Ω. Then*

$$\pi(m_1, m_2) \in q_{ij} \quad \textit{if and only if} \quad \pi(m_1, m_1) = \delta_i \wedge \pi(m_2, m_2) = \delta_j.$$

Proof. Observe that $\pi(m_1, m_2) \in q_{ij}$ implies that $\pi(m_1, m_2)$ and δ_i are E-related. Therefore there must be some m_2' such that $\pi(m_1, m_2') = \delta_i$. As π respects Δ, $m_2' = m_1$, so that $\pi(m_1, m_1) = \delta_i$. Applying the same argument to E^T we get $\pi(m_2, m_2) = \delta_j$.

Conversely, $\pi(m_1, m_1) = \delta_i \in q_i$ implies $\pi(m_1, m_2) \in q_i$ for all m_2. Repeating the same argument we get that $\pi(m_2, m_2) = \delta_j$ implies $\pi(m_2, m_1) \in q_j$ so that $\pi(m_1, m_2) \in q_j^T$. Putting these together, $\pi(m_1, m_1) = \delta_i$ and $\pi(m_2, m_2) = \delta_j$ imply $\pi(m_1, m_2) \in q_{ij}$. □

We define the restrictions of a game tableau Ω to its subdomains q_{ij}. Note that in restriction to each q_{ij} the equivalence relations E and E^T become trivial since the q_{ij} are the classes of the common refinement of the two. As for Δ it is obvious that $\Delta \cap q_{ij} = \emptyset$ if $i \neq j$, and $\Delta \cap q_{ii} = \{\delta_i\}$. T is an involutive mapping from q_{ii} to itself, and turns into a bijection between q_{ij} and q_{ji} for $i \neq j$. If we also consider weighted tableaux, it makes sense to retain both weight functions, ν and ν^T over each q_{ij} with $i \neq j$ since T is no longer internal to q_{ij}.

Definition 7.17. *Let $\mathfrak{Q}$ be a game tableau, the q_{ij} as defined in Definition 7.15. The restriction of $\mathfrak{Q}$ to q_{ij} is defined to be*

$$\mathfrak{Q}_{ij} := \begin{cases} \big(q_{ii}, \leqslant\restriction q_{ii}, T \restriction q_{ii}, \{\delta_i\}\big) & \text{for the diagonal case } j = i \\ \big(q_{ij}, \leqslant\restriction q_{ii}\big) & \text{for the off-diagonal case } i \neq j. \end{cases}$$

For the restrictions of a weighted tableau $(\mathfrak{Q}; \nu)$, put

$$(\mathfrak{Q}; \nu)_{ij} := \begin{cases} \big(\mathfrak{Q}_{ii}; \nu \restriction q_{ii}\big) & \text{for } j = i \\ \big(\mathfrak{Q}_{ij}; \nu \restriction q_{ij}, \nu^T \restriction q_{ij}\big) & \text{for } i \neq j. \end{cases}$$

This decomposition calls for an adapted notion of realizations. The modifications and simplifications required with respect to Definition 7.7 are canonical. The diagonal restrictions to the q_{ii} can in fact be regarded as special cases of game tableaux, with trivial E. It is only for the off-diagonal boxes that formal modifications are required.

Definition 7.18. *Let $\mathfrak{Q}$ be a game tableau, $\mathfrak{Q}_{ij}$ its restriction to a subdomain q_{ij}. Assume first $i \neq j$. A surjective mapping $\pi: s \times t \to q_{ij}$ is a* realization *of $\mathfrak{Q}_{ij}$, if for all $m_1 \in s$ and all $m_2 \in t$:*

$$\big\{\pi(m_1, m_2') \bigm| m_2' \in t\big\} = \big\{\pi(m_1', m_2) \bigm| m_1' \in s\big\} = q_{ij}.$$

π realizes the weighted restriction $(\mathfrak{Q}; \nu)_{ij} = (\mathfrak{Q}_{ij}; \nu \restriction q_{ij}, \nu^T \restriction q_{ij})$ if for all $\alpha \in q_{ij}$, $m_1 \in s$ and all $m_2 \in t$:

$$\begin{aligned} \big|\big\{m_2' \in t \bigm| \pi(m_1, m_2') = \alpha\big\}\big| &= \nu(\alpha), \\ \big|\big\{m_1' \in s \bigm| \pi(m_1', m_2) = \alpha\big\}\big| &= \nu^T(\alpha). \end{aligned}$$

For $i = j$ the conditions that a surjective mapping $\pi: s \times s \to q_{ii}$ realizes $\mathfrak{Q}_{ii}$ or $(\mathfrak{Q}, \nu)_{ii}$ are those of Definition 7.7 applied to the game tableau $\big(q_{ii}, \leqslant\restriction q_{ii}, q_{ii} \times q_{ii}, T \restriction q_{ii}, \{\delta_i\}\big)$ and to the weighted game tableau $\big(q_{ii}, \leqslant\restriction q_{ii}, q_{ii} \times q_{ii}, T \restriction q_{ii}, \{\delta_i\}; \nu \restriction q_{ii}\big)$.

Suppose $\pi: n \times n \to Q$ is a realization of $\mathfrak{Q}$ or $(\mathfrak{Q}, \nu)$ over n. Let n be decomposed into the subsets $\big\{m \in n \bigm| \pi(m, m) = \delta_i\big\}$. Obviously, n is the disjoint union of these. Without loss of generality we may assume that n is the disjoint ordered sum $n = \sum_i n_i$, $n_i = \big\{m \in n \bigm| \pi(m, m) = \delta_i\big\}$. Formally this means that n_1 is identified with an initial subset of n, n_2 consists of a consecutive interval following that initial segment and so on. In particular *identifying n with the disjoint ordered sum* of the n_i implies that the subsets n_i are embedded in a well-defined way into n such that n is the disjoint union of the embedded n_i. In the present case this situation may be assumed without loss of generality because a realization over n can be composed with any permutation of n in the obvious manner to yield a new realization, equivalent with the former one for our purposes.

With such a presentation of $n = \sum_i n_i$, $n_i = \{m \in n \mid \pi(m,m) = \delta_i\}$, we immediately have that the restrictions of π to the subsets $n_i \times n_j \subseteq n \times n$ provide realizations for the restrictions Ω_{ij} and $(\Omega;\nu)_{ij}$. It is straightforward to check the conditions mentioned in the last definition. Let us supply the argument for surjectivity of $\pi_{ij} := \pi \restriction n_i \times n_j \to q_{ij}$: by Lemma 7.16, $\pi(m_1,m_2) \in q_{ij}$ if and only if $\pi(m_1,m_1) = \delta_i$ and $\pi(m_2,m_2) = \delta_j$, i.e. if and only if $m_1 \in n_i$ and $m_2 \in n_j$. Thus surjectivity of π_{ij} follows from surjectivity of π itself.

The interesting fact is that, conversely, realizations of the individual restrictions can be fit together to form a realization of the whole (weighted) tableau if they satisfy just the most obvious compatibility conditions relating the sizes of the subdomains.

Lemma 7.19. *Let Ω be a game tableau, Ω_{ij}, $1 \leqslant i,j \leqslant l$, its restrictions to the q_{ij} defined as above. Assume that for some tuple $(n_i)_{1\leqslant i\leqslant l}$ of positive numbers there are surjective mappings $\pi_{ij}: n_i \times n_j \to q_{ij}$ for each $1 \leqslant i \leqslant j \leqslant l$, such that π_{ij} is a realization of Ω_{ij}. Then the following is a realization of Ω on the disjoint ordered sum $n := \sum_i n_i$:*

$$\pi(m_1,m_2) := \begin{cases} \pi_{ij}(m_1,m_2) & \text{if } m_1 \in n_i,\ m_2 \in n_j,\ i \leqslant j, \\ T\pi_{ji}(m_2,m_1) & \text{if } m_1 \in n_i,\ m_2 \in n_j,\ j < i. \end{cases}$$

The same holds for realizations of a weighted tableau $(\Omega;\nu)$.

Before giving a proof, let us note that together with the preceding considerations we have thus found that a (weighted) game tableau is realizable if and only if its restrictions are realizable over subdomains of matching sizes. In terms of at first arbitrary domains for the realization of the restrictions, $\pi_{ij}: n_1^{ij} \times n_2^{ij} \to q_{ij}$ the conditions for *matching size* are that the n_1^{ij} are independent of j, and that the n_2^{ij} are independent of i.

Proposition 7.20. *A game tableau Ω has a realization over n if and only if $n = \sum_i n_i$ for $n_i > 0$ such that each of its restrictions Ω_{ij} admits a realization over $n_i \times n_j$.*

Similarly, a weighted game tableau $(\Omega;\nu)$ has a realization over n if and only if $n = \sum_i n_i$ such that its restrictions $(\Omega;\nu)_{ij}$ admit realizations over $n_i \times n_j$, for all i,j.

Proof (of Lemma 7.19). Recall from Definition 7.7 the conditions on a realization π of Ω. It is clear that π as defined above is surjective, since the π_{ij} are surjective mappings to the q_{ij} and since $Q = \bigcup q_{ij} = \bigcup_{i \leqslant j} q_{ij} \cup \bigcup_{j<i} Tq_{ji}$ as $q_{ij} = Tq_{ji}$. π respects the diagonal, because the π_{ii} do. π commutes with T, because the π_{ii} do and because the appropriate transformation under T is explicitly built into π on the off-diagonal boxes. We check that π is correct with respect to E. By construction, $\pi(m_1, m_2) \in q_{ij}$ if and and only if $m_1 \in n_i$ and $m_2 \in n_j$. Varying $m_2' \in n_{j'}$ we get $\{\pi(m_1, m_2') \mid m_2' \in n_{j'}\} = q_{ij'}$, since the corresponding behaviour is required of the $\pi_{ij'}$ or $\pi_{j'i}$. Therefore $\{\pi(m_1, m_2') \mid m_2' \in n\} = q_i$ as required.

The proof for a weighted tableau is similar. The multiplicity requirements for the realizations π_{ij} of the restrictions immediately imply that also the composition π realizes the multiplicities prescribed by the overall weight function. □

We now pursue the actual constructions of realizations in separate presentations for $L^2_{\infty\omega}$ and $C^2_{\infty\omega}$. There are more constraints in the case of $C^2_{\infty\omega}$, so that the constructions are more difficult. On the other hand these constructions appear more straightforward because there are more data available and correspondingly fewer arbitrary choices to be made. We treat $C^2_{\infty\omega}$ or the realization of weighted tableaux first and specialize and modify this treatment in Section 7.3 to obtain the results for $L^2_{\infty\omega}$.

7.2 Realizations for I_{C^2}

7.2.1 Necessary Conditions

The numerical information contained in the I_{C^2}, and in the weighted game tableaux that derive from these, fixes the size of a possible realization and the n_i as in the compatibility conditions in Proposition 7.20. Fix a weighted game tableau $(\Omega; \nu)$, with q_i, q_{ij}, δ_i for $1 \leqslant i, j \leqslant l$ defined according to Definition 7.15. Further define the numbers

$$n_j := \sum_{\alpha \in q_{ij}} \nu(\alpha) \quad \text{and} \quad n := \sum_i n_i.$$

An equivalent definition of n_i in terms of $(\Omega; \nu)_{ij}$ is $n_i = \sum_{\alpha \in q_{ij}} \nu^T(\alpha)$. Equivalence with the above is a consequence of the fact that T is an involutive bijection between q_{ij} and q_{ji} and that, by definition, $\nu^T = \nu \circ T$.

Lemma 7.21. *Let the n_i and n be as just defined. If $(\Omega; \nu)$ has any realization, then it must be over n. The induced realizations of the restrictions $(\Omega; \nu)_{ij}$ must be over domains $n_i \times n_j$.*

Proof. Let π be a realization of $(\Omega; \nu)$ over s. By Lemma 7.16, $\pi(m_1, m_2) \in q_{ij}$ if and only if $\pi(m_1, m_1) = \delta_i$ and $\pi(m_2, m_2) = \delta_j$. Choose $m_1 \in s$ such that $\pi(m_1, m_1) = \delta_i$. Then

$$\left|\{m_2 \in s \mid \pi(m_1,m_2) \in q_{ij}\}\right| \;=\; \textstyle\sum_{\alpha\in q_{ij}} \nu(\alpha) \;=\; n_j\,.$$

It further follows that also $n_j = \left|\{m_2 \in s \mid \pi(m_2,m_2) = \delta_j\}\right|$.

Note in particular, that these numbers depend on j and not on i. This is just the compatibility condition of Proposition 7.20. Applying the same argument to variations in the first component, and with fixed m_2 for which $\pi(m_2,m_2) = \delta_j$, we obtain

$$\begin{aligned}\left|\{m_1 \in s \mid \pi(m_1,m_2) \in q_{ij}\}\right| &= \textstyle\sum_{\alpha\in q_{ij}} \nu^T(\alpha) \\ &= \textstyle\sum_{\alpha\in q_{ji}} \nu(\alpha) \;=\; n_i.\end{aligned}$$

This shows that the induced realization of $(\Omega;\nu)_{ij}$ must be over $n_i \times n_j$. The first claim of the lemma, $s = \sum_i n_i$, follows directly from Proposition 7.20. □

The following lemma states some necessary conditions for the realizability of $(\Omega;\nu)$ in terms of the restrictions $(\Omega;\nu)_{ij}$. Sufficiency of these conditions will be shown in the sequel.

Lemma 7.22. *Any realization of $(\Omega;\nu)_{ij}$ is over $n_i \times n_j$. Recall that the numbers n_i and n_j are defined in terms of $(\Omega;\nu)_{ij}$ as*

$$n_i = \sum_{\alpha\in q_{ij}} \nu^T(\alpha) \quad \textit{and} \quad n_j = \sum_{\alpha\in q_{ij}} \nu(\alpha).$$

If $(\Omega;\nu)_{ij}$ has a realization then for all $\alpha \in q_{ij}$:

$$(*) \qquad \frac{\nu(\alpha)}{\nu^T(\alpha)} = \frac{n_j}{n_i}.$$

For realizability of a diagonal restriction $(\Omega;\nu)_{ii}$ it is necessary that in addition $\nu(\delta_i) = 1$, and that if n_i is odd, then for all $\alpha \in q_{ii} \setminus \{\delta_i\}$:

$$(**) \qquad T\alpha = \alpha \implies \nu(\alpha) \textit{ is even.}$$

Proof. Suppose that $\pi\colon s\times t \to q_{ij}$ realizes $(\Omega;\nu)_{ij}$ (cf. Definition 7.18). That $s \times t = n_i \times n_j$ is shown by an argument similar that in Lemma 7.21, but in restriction to the individual q_{ij}. We show that $t = n_j$. For all $m_1 \in s$ and all $\alpha \in q_{ij}$, $\left|\{m_2 \in t \mid \pi(m_1,m_2) = \alpha\}\right| = \nu(\alpha)$. Therefore $t = \left|\{m_2 \in t \mid \pi(m_1,m_2) \in q_{ij}\}\right| = \sum_{\alpha\in q_{ij}} \nu(\alpha) = n_j$.

For the quotient conditions $(*)$ it suffices to count the number of pairs that are mapped to α, first in column-wise fashion, then row-wise and equate the two:

$$\begin{aligned}\left|\{(m_1,m_2) \mid \pi(m_1,m_2) = \alpha\}\right| &= \textstyle\sum_{m_1\in s}\left|\{m_2 \mid \pi(m_1,m_2) = \alpha\}\right| \\ &= \textstyle\sum_{m_1\in s} \nu(\alpha) \;=\; s\,\nu(\alpha),\end{aligned}$$

$$\begin{aligned}\left|\{(m_1,m_2) \mid \pi(m_1,m_2) = \alpha\}\right| &= \textstyle\sum_{m_2\in t}\left|\{m_1 \mid \pi(m_1,m_2) = \alpha\}\right| \\ &= \textstyle\sum_{m_2\in t} \nu^T(\alpha) \;=\; t\,\nu^T(\alpha).\end{aligned}$$

Consider now the additional constraints expressed for the diagonal case. Necessity of $\nu(\delta_i) = 1$ is obvious, since π must respect $\Delta \restriction q_{ii} = \{\delta_i\}$. For (**) assume $s = t = n_i$ is odd and that $\alpha \neq \delta_i$ is a fixed point under T. $\pi^{-1}(\alpha)$ must be disjoint from the diagonal $\{(m,m) | m \in s\}$, because $\pi(m,m) = \delta_i$. This implies that T operates as a fixed-point free involutive bijection on $\pi^{-1}(\alpha)$. Therefore $|\pi^{-1}(\alpha)|$ must be even. The above counting equations imply that this number equals $|\{(m_1, m_2) \mid \pi(m_1, m_2) = \alpha\}| = s\,\nu(\alpha)$. If s is odd, therefore, $\nu(\alpha)$ must be even. □

7.2.2 Realizations of the Off-Diagonal Boxes

We turn to the proof of sufficiency of the conditions expressed in the last lemma. The realization of off-diagonal restrictions turns out to be quite straightforward.

Lemma 7.23. *Let $(\Omega; \nu)_{ij}$, $i \neq j$ satisfy condition (*) of Lemma 7.22: for all α $\nu(\alpha)/\nu^T(\alpha) = n_j/n_i$. Then there is a realization $\pi: n_i \times n_j \to q_{ij}$ of $(\Omega; \nu)_{ij}$. Such realizations are constructible in time polynomial in $n_i n_j$.*

Proof. Let t/s be the reduced presentation of n_j/n_i. Let $\alpha_1, \ldots, \alpha_r$ be an enumeration of q_{ij} as ordered by $\leqslant$. By assumption there exist numbers d_k, for $1 \leqslant k \leqslant r$, such that $\nu(\alpha_k) = d_k t$ and $\nu^T(\alpha_k) = d_k s$. Putting $d = \sum d_k$ we have $ds = n_i$ and $dt = n_j$ by definition of n_i and n_j. For the following compare Figure 7.2.

Fig. 7.2

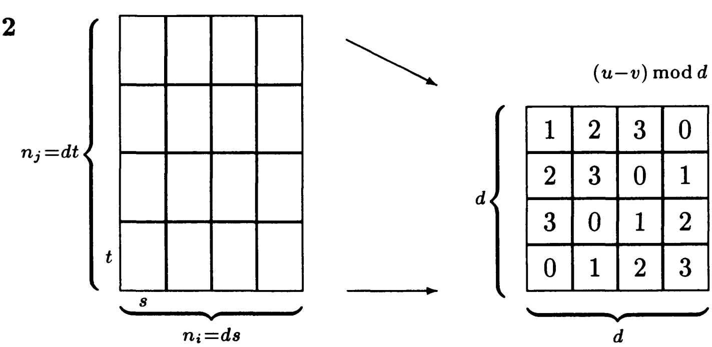

Identify d with the ordered disjoint sum of the d_k and let $g: d \to \{1, \ldots, r\}$ be the function that characterizes the embedded d_k: $g(u)$ is that k with $u \in d_k$. Also identify n_i with the product $d \times s$, and similarly n_j with $d \times t$. Note that the sum and product identifications can be uniquely defined with the help of the natural orderings (we have done this explicitly for the sum above). Define a surjective function

$$\begin{aligned} f\colon (d \times s) \times (d \times t) &\longrightarrow m \\ ((u,x),(v,y)) &\longmapsto (u-v) \bmod d\,. \end{aligned}$$

Figure 7.2 sketches the situation in an example with $d = 4$. The passage from the left to the right indicates the effect of the projections $n_i = d \times s \to d$ and $n_j = d \times t \to d$ involved in the definition of f. On the right-hand side the distribution of values for the function $(u-v) \bmod d$ is indicated.

A realization π can now be defined on $n_i \times n_j = (d \times s) \times (d \times t)$ by

$$\pi(m_1, m_2) := \alpha_{g(f(m_1,m_2))}.$$

π factorizes with respect to f and maps all those blocks, whose values under f fall into $d_k \subseteq d$, to α_k. Let us check that π realizes the multiplicities for columns as specified by $\nu \restriction q_{ij}$. Consider $m_1 = (u,x) \in d \times s$:

$$\big\{m_2 \in d \times t \;\big|\; \pi(m_1,m_2) = \alpha_k\big\} = \big\{m_2 \in d \times t \;\big|\; f(m_1,m_2) \in d_k\big\}.$$

$f(m_1,(v,y)) = (u-v) \bmod d$ so that there are exactly $d_k t = \nu(\alpha_k)$ many $m_2 = (v,y)$ such that $f(m_1,m_2) \in d_k$. The multiplicity conditions on rows are checked to be in accordance with $\nu^T \restriction q_{ij}$ in exactly the same way.

Note that the proposed construction of a realization is quite definite: we have sketched how to construct a particular solution to the realization problem for an off-diagonal restriction. This construction is clearly in PTIME with respect to the product $n_i n_j$. □

7.2.3 Magic Squares

Sufficiency of the conditions of Lemma 7.22 and the construction of realizations is combinatorially more demanding for diagonal restrictions $(\Omega;\nu)_{ii}$ because of the symmetries imposed by T. We first present a preparatory lemma on certain colourings of squares.

The most complicated case in the construction of realizations for diagonal restrictions $(\Omega;\nu)_{ii}$ — the case of even n_i with fixed points $\alpha \neq \delta_i$ under T, as we shall see — reduces to the construction of such colourings. The colourings described in the following lemma in fact present the worst case for the construction of a realization. The symmetry requirements for these colourings are reminiscent of magic squares and related number puzzles (and call for a try with paper and pencil).

Lemma 7.24. *Let n be even. Then there is a colouring $c\colon n \times n \longrightarrow n$ of the n-square $n \times n$ with n colours $0,\ldots,n-1$ with the following properties:*

(i) the main diagonal, i.e. all identity pairs, are coloured 0.
(ii) each colour occurs exactly once in each row and in each column.
(iii) the entire colouring is mirror symmetric with respect to the main diagonal. In other words the colouring is invariant under T.

A colouring of this kind can be constructed in time polynomial in n.

Observe that the same puzzle cannot be solved for odd n: each colour apart from 0 has to occur an even number of times because T operates as a fixed-point free involutive mapping on the points of this colour.

An example of a colouring of the 6×6 square according to the requirements of the lemma is given in Figure 7.5 (a) below.

Proof. We give an inductive existence proof that can immediately be turned into a PTIME construction. The claim is obvious for $n = 2$ and we now show how to construct good colourings of the $2n$-square and the $(2n-2)$-square from a given good colouring of the n-square. This yields a valid inductive proof, because for even $m > 2$ at least one of $m/2$ and $(m+2)/2$ is even and smaller than m. Let $c: n \times n \to n$ be a good colouring.

Fig. 7.3

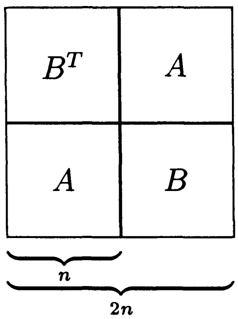

A good colouring $C: (2n) \times (2n) \to 2n$ is easily obtained by gluing four copies of trivially modified c-coloured squares together. The pattern is indicated in Figure 7.3. The box A represents an n-square coloured according to c, B an n-square coloured by $c': n \times n \to \{n, \ldots, 2n-1\}$, $c'(m_1, m_2) = c(m_1, m_2) + n$. B^T finally is coloured $c' \circ T$.

Fig. 7.4

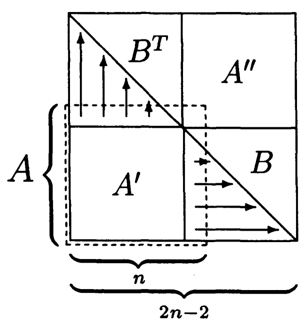

Consider now the $(2n-2)$-square. Assume without loss of generality that the c-coloured n-square A has top row (from left to right) coloured $n-1$, $n-2$, …, 0. By symmetry of c this implies that the rightmost column of A

is coloured $0, 1, \ldots, n-1$ (from top to bottom). Let A' be the $(n-1)$-square coloured by $c' = c \restriction (n-1) \times (n-1)$, or A with top row and rightmost column removed. Let A'' be the mirror image of A' across the second diagonal. As a second building block we use an $(n-1)$-square B coloured as follows. The second diagonal of B is coloured $1, \ldots, n-1$ from top left to bottom right. The remaining places are coloured with colours $n, \ldots, 2n-3$ such that each of these colours occurs exactly once in each row and in each column. This can be done with a cyclic permutation of colours following the second diagonal. B^T finally is the mirror image of B across the main diagonal. A good colouring of the $(2n-2)$-square is obtained by gluing the four $(n-1)$-squares squares A', A'', B and B^T as indicated in Figure 7.4. Note that the second diagonals in B and B^T exactly supply those colours from n in each row and in each column, that are missing in A' and A''. The arrows in the figure indicate how these second diagonals in B and B^T replace the rows and columns cut away from A. □

Lemma 7.25. *Let $D \subseteq n \times n$ be a subset of the n-square that is symmetric with respect to the main diagonal (invariant under T), disjoint from the main diagonal, and contains exactly two elements of each row and of each column. Then there is a colouring of this subset with two colours $c\colon D \to \{0,1\}$ such that each colour occurs exactly once in each row and in each column and such that c is antisymmetric with respect to the main diagonal: $T \circ c = 1 - c$. Such c is* PTIME *computable from D.*

In Figure 7.5 (b) the set D consisting of those points that are coloured 2 or 5 in (a) is split according to these requirements.

Proof. Consider the relation S of belonging to the same row *or* to the same column of $n \times n$ in restriction to D. Since each row and each column contains exactly 2 elements of D, D must be the disjoint union of even-length S-cycles. Since S and D are T-invariant, it follows that for each such cycle C either $T(C) \cap C = \emptyset$ or $T(C) = C$. Two different cycles cannot contain points of the same row or of the same column, since for instance there are only two points of D in any row and these necessarily belong to the same cycle. It follows that the requirements on c can be satisfied if for each cycle C there is a colouring $c'\colon C \cup T(C) \to \{0,1\}$, which is antisymmetric for T and contains at most one point coloured 0, respectively 1, in each row and each column. To obtain c take the union of these c'. Consider first a single cycle C. C can be coloured alternately with colours 0 and 1, starting with colour 0 say from the lexicographically least member in C and proceeding in the direction of the horizontal S-neighbour of this point. If $T(C) \neq C$, then c' on $C \cup T(C)$ can be taken as the union of this colouring of C with the antisymmetric image under T on $T(C)$. In case $T(C) = C$ we have to check that the colouring we have obtained is antisymmetric itself. Let C be enumerated in the order used in the colouring procedure as $c_0, c_1, \ldots, c_{2n} = c_0$, so that c_0 and c_1 are in the same row, c_1 and c_2 in the same column, etc. Assume for contradiction that for example $Tc_{2i} = c_{2j}$ (c_{2i} and c_{2j} are both coloured 0), where $i < j$. Then c_{2i+1}, the element of D in the same row as c_{2i}, must be T-related with c_{2j-1}, which is the element of D that is in the same column with c_{2j}. This is because D is symmetric with respect to T. Proceeding in this manner we would find $Tc_k = c_k$ for $k = i + j$, which is impossible since D is disjoint from the diagonal. □

Fig. 7.5 **(a)**

3	5	4	2	1	0
4	2	5	3	0	1
5	4	1	0	3	2
2	3	0	1	5	4
1	0	3	4	2	5
0	1	2	5	4	3

(b)

	●		○		
	○	●			
○					●
●				○	
				●	○
		○	●		

7.2.4 Realizations of the Diagonal Boxes

With the help of these lemmas about colourings we finally construct realizations for diagonal restrictions $(\Omega; \nu)_{ii}$ that satisfy the conditions derived in Lemma 7.22.

Lemma 7.26. *Let $(\Omega; \nu)_{ii}$ satisfy the following conditions:*

(i) $\nu(\delta_i) = 1$.
(ii) for all $\alpha \in q_{ii}$: $\nu(\alpha) = \nu^T(\alpha)$.
(iii) if $n_i = \sum_{\alpha \in q_{ii}} \nu(\alpha)$ is odd, then $\nu(\alpha)$ is even for all $\alpha \neq \delta_i$ with $T\alpha = \alpha$.

Then there is a realization $\pi: n_i \times n_i \to q_{ii}$ of $(\Omega; \nu)_{ii}$ over $n_i = \sum_{\alpha \in q_{ii}} \nu(\alpha)$. Moreover, such a realization can be constructed in time polynomial in n_i.

Proof. Recall from Definitions 7.18 and 7.7 that a realization π of $(\Omega; \nu)_{ii}$ has to satisfy the following:

(a) $\pi(m_1, m_2) = \delta_i$ exactly for $m_1 = m_2$.
(b) $T \circ \pi = \pi \circ T$.
(c) For all m_1: $\left|\{m_2 \mid \pi(m_1, m_2) = \alpha\}\right| = \nu(\alpha)$.

(c) is a combination of (iii) and (iv) in Definition 7.7 applied to the present case with trivial E. Note that (c) together with (b) also implies that for all m_2: $\left|\{m_1 \mid \pi(m_1, m_2) = \alpha\}\right| = \left|\{m_1 \mid \pi(m_2, m_1) = T\alpha\}\right| = \nu(T\alpha)$, which by assumption (ii) of the lemma is the same as $\nu(\alpha)$.

The construction of π depends on whether n_i is even or odd. The odd case is the easier one.

Case A: n_i odd. Let $q_{ii} \setminus \{\delta_i\} = q_0 \dot\cup q_1$ where q_0 consists of those points that are fixed under T. Since T is a fixed-point free involutive permutation on q_1, $q_1 = q \dot\cup Tq$ for some $q \subseteq q_1$. For a definite construction it is important that q can be specified in a unique way, with the help of the ordering $\leqslant$ on q_{ii}. For instance we may take q as the lexicographically least subset $q \subseteq q_1$ for which $q_1 = q \dot\cup Tq$. This q can be determined in PTIME. Let now q_0 be enumerated as $\alpha_1, \dots, \alpha_s$, and q as $\beta_1, \dots, \beta_t$, both in $\leqslant$-order. So $\delta_i, \alpha_1, \dots, \alpha_s, \beta_1, \dots, \beta_t, T\beta_1, \dots, T\beta_t$ is an enumeration of q_{ii} without repetitions. By the assumptions of the lemma, $\nu(\alpha_j) = 2k_j$ for suitable k_j, $\nu(\beta_j) = \nu(T\beta_j) =: l_j$. Let $d_1 := \sum k_j$, $d_2 := \sum l_j$, $d := d_1 + d_2$, where we take these as identifications with the disjoint ordered sums. Note that $n_i = 1 + 2d$. Put $D_0 := \{(u,v) \in n_i \times n_i \mid (u - v) \bmod n_i \in \{1, \dots, d\}\}$. It follows that $n_i \times n_i = \{(u,u) \mid u \in n_i\} \dot\cup D_0 \dot\cup T(D_0)$. See the sketch in Figure 7.6 with $d = 3$, $n_i = 7$, where the values $(u - v) \bmod n_i$ in $\{1, \dots, d\}$ are indicated. The desired realization π can be defined as follows:

$$\pi(u,v) := \begin{cases} \delta_i & \text{for } u = v \\ \alpha_j & \text{if } (u,v) \in D_0 \text{ and } (u-v) \bmod n_i \in k_j \\ \beta_j & \text{if } (u,v) \in D_0 \text{ and } (u-v) \bmod n_i \in l_j \\ T(\pi(v,u)) & \text{if } (u,v) \in T(D_0)\,. \end{cases}$$

Conditions (a) and (b) are obviously satisfied.

To check the multiplicity requirements (c), note that for each $u \in n_i$ and $s \in \{1, \dots, d\}$, $\left|\{v | (u-v) \bmod n_i = s\}\right| = \left|\{v | (u-v) \bmod n_i = -s\}\right|$ and that the operation of T on (u,v) translates into $(u-v) \bmod n_i \mapsto -(u-v) \bmod n_i$. It follows that for all $u \in n_i$ indeed $\left|\{v | \pi(u,v) = \alpha_j\}\right| = 2k_j = \nu(\alpha_j)$ and $\left|\{v | \pi(u,v) = \beta_j\}\right| = l_j = \nu(\beta_j)$ as required.

Fig. 7.6

1	2	3				
2	3					1
3					1	2
				1	2	3
			1	2	3	
		1	2	3		
	1	2	3			

Case B: n_i even. Enumerate q_{ii} as $\delta_i, \alpha_1, \dots, \alpha_s, \beta_1, \dots, \beta_t, T\beta_1, \dots, T\beta_t$ as above. In particular $T\alpha_j = \alpha_j$. Let $\nu(\alpha_j) = k_j$, $\nu(\beta_j) = \nu(T\beta_j) = l_j$. Thus $n_i = 1 + d_1 + 2d_2$, where $d_1 = \sum k_j$, $d_2 = \sum l_j$. We identify n_i with the disjoint ordered sum $1 + \sum k_j + 2d_2$.

Let $c\colon n_i \times n_i \to n_i$ be a colouring function as constructed in Lemma 7.24. The crucial properties are symmetry, $c \circ T = c$, exactly one occurrence of every colour in each row and in each column, and $c(u,u) = 0$ on the diagonal. With the above identification we consider c as a function to the disjoint sum $1 + \sum k_j + 2d_2$. Thus c can directly be used to define π partially as

$$\pi(u,v) := \begin{cases} \delta_i & \text{for } u = v \\ \alpha_j & \text{if } c(u,v) \in k_j \,. \end{cases}$$

Symmetry and multiplicities for the α_j are as required. It remains to define $\pi \restriction \{(u,v) \mid c(u,v) \in 2d_2\}$ with values in $\{\beta_1, \ldots, \beta_t, T\beta_1, \ldots, T\beta_t\}$. Note that this remaining subdomain is disjoint from the diagonal. We now further identify $2d_2$ with the product $\{0,1\} \times \sum l_j$ (in some canonical and definite way), so that on the remaining domain c takes values $(0,d)$ and $(1,d)$ for $d \in \dot{\bigcup} l_j$. Put $D_d := \{(u,v) \mid c(u,v) = (0,d) \vee c(u,v) = (1,d)\}$ for $d \in \dot{\bigcup} l_j$. D_d is T-symmetric, disjoint from the diagonal and contains exactly two elements of each row and each column. By Lemma 7.25 each D_d can be coloured by some $c_d \colon D_d \to \{0,1\}$ in such a way that each column and each row contains colour 0 and 1 exactly once, and such that $T \circ c_d = 1 - c_d$ corresponds to an inversion of the colouring. To complete the definition of π put

$$\pi(u,v) := \begin{cases} \beta_j & \text{if } (u,v) \in D_d, d \in l_j \text{ and } c_d(u,v) = 0 \\ T\beta_j & \text{if } (u,v) \in D_d, d \in l_j \text{ and } c_d(u,v) = 1 \,. \end{cases}$$

Compatibility with T follows, since T preserves D_d and inverts c_d. The multiplicities for the β_j are realized correctly because each row and each column contains exactly one element $(u,v) \in D_d$ such that $c_d(u,v) = 0$ (respectively 1) for each d. Therefore $|\{v|\pi(u,v) = \beta_j\}| = |\{v|\pi(u,v) = T\beta_j\}| = l_j$ as required. □

Putting those results of the preceding sections, that relate to the case of $C^2_{\infty\omega}$, together we have the following.

Proposition 7.27. *Let $(\Omega;\nu)$ be a weighted game tableau. Let the E-classes of Q be $q_1, \ldots, q_l$, let $q_{ij} := q_i \cap T(q_j)$ and put $n_i := \sum_{\alpha \in q_{ii}} \nu(\alpha)$. $(\Omega;\nu)$ admits a realization if and only if the following conditions are satisfied:*

(i) $\sum_{\alpha \in q_{ij}} \nu(\alpha) = n_j$ *independent of* i.
(ii) $\nu(\delta) = 1$ *for all* $\delta \in \Delta$.
(iii) $\nu(\alpha) \,/\, \nu^T(\alpha) = n_j \,/\, n_i$ *for all* $\alpha \in q_{ij}$.
(iv) For all odd n_i, *and all* $\alpha \in q_{ii} \setminus \Delta$, *if* $T\alpha = \alpha$, *then* $\nu(\alpha)$ *is even.*

In this case a realization on $n = \sum n_i$ can be constructed in time polynomial in n, thus proving the $C^2_{\infty\omega}$-related part of Theorem 7.14.

We review the arguments that lead to this statement: (i) is the compatibility condition for fitting together realizations of the restrictions $(\Omega;\nu)_{ij}$; necessity follows from Proposition 7.20 together with Lemma 7.22. (ii) is obviously necessary. (iii) is necessary for realizability of each $(\Omega;\nu)_{ij}$, (iv) is necessary for $(\mathfrak{Q};\nu)_{ii}$ to admit a realization, both by Lemma 7.22. Sufficiency follows from realizability of the $(\mathfrak{Q};\nu)_{ij}$: (iii) suffices for $(\mathfrak{Q};\nu)_{ij}$, $i \neq j$, see Lemma 7.23; (ii) – (iv) suffice for $(\mathfrak{Q};\nu)_{ii}$ according to Lemma 7.26; (i) suffices to compose these individual realizations.

7.3 Realizations for I_{L^2}

7.3.1 Necessary and Sufficient Conditions

We prove the following analogue of Proposition 7.27 in the case of game tableaux without weights.

Proposition 7.28. *Let $\mathfrak{Q}$ be a game tableau. Let the E-classes of Q be $q_1, \dots, q_l$. Put $q_{ij} := q_i \cap T(q_j)$. Then $\mathfrak{Q}$ admits a realization if and only if the following conditions are satisfied:*

(i) all q_{ij} are nonempty.
(ii) if $|q_{ii}| = 1$ then $|q_{ij}| = 1$ for all j.

In this case a realization — one of minimal size even — can be constructed in time polynomial in $|Q|$. This proves that part of Theorem 7.14 that relates to $L^2_{\infty\omega}$.

Proof (of necessity of (i) and (ii)). (i) is trivial: if $\mathfrak{Q}_{ij}$ is to have a realization, then q_{ij} must not be empty. For (ii) assume that $|q_{ii}| = 1$, i.e. that $q_{ii} = \{\delta_i\}$. It follows that $\mathfrak{Q}_{ii}$ can only admit the trivial realization $\pi: 1 \times 1 \to \{\delta_i\}$ on the one-element square, since no off-diagonal pair may be mapped to δ_i by any realization. In the terminology of Proposition 7.20 it follows that $n_i = 1$ and that all $\mathfrak{Q}_{ij}$ must have realizations on domains $1 \times n_j$. It follows directly from Definition 7.18 that any realization $\pi: s \times t \to q_{ij}$ satisfies $s, t \geqslant |q_{ij}|$, so that $s = 1$ implies $|q_{ij}| = 1$. □

The rest of this section is devoted to the proof of the sufficiency claim of Proposition 7.28. Again Proposition 7.20 is invoked to reduce the construction of a realization for $\mathfrak{Q}$ to the realization of the restrictions $\mathfrak{Q}_{ij}$. In fact we shall see that (ii) in the proposition reflects what remains of the numerical compatibility conditions in Proposition 7.20 in the case of $L^2_{\infty\omega}$: $L^2_{\infty\omega}$ can only count "0, 1, many". For the restrictions, we first treat the off-diagonal ones, then the diagonal ones. Fix a game tableau $\mathfrak{Q}$.

The off-diagonal restrictions.

Lemma 7.29. *For $i \neq j$. If $q_{ij} \neq \emptyset$ then there are realizations $\pi: s \times t \to q_{ij}$ of $\mathfrak{Q}_{ij}$ exactly for all $s, t \geqslant |q_{ij}|$.*

Sketch of Proof. The condition $s, t \geqslant |q_{ij}|$ is necessary since π has to attain each $\alpha \in q_{ij}$ at least once in every row and in every column.

Let $s_0 = |q_{ij}|$ and first construct a realization on $s_0 \times s_0$. Let q_{ij} be enumerated as $\alpha_0, \ldots \alpha_{s_0-1}$ in increasing order with respect to $\leqslant$. Put

$$\pi_0(u,v) := \alpha_k \qquad \text{for } k = (u+v) \bmod s_0.$$

Obviously each α_j occurs once in each row and in each column as required. To obtain realizations for $s \times t$, $s, t \geqslant s_0$ put

$$\pi(u,v) := \begin{cases} \pi_0(u,v) & \text{for } u, v < s_0 \\ \alpha_k & \text{for } k = v \bmod s_0, u \geqslant s_0, v \leqslant s_0 \\ \alpha_k & \text{for } k = u \bmod s_0, v \geqslant s_0, u \leqslant s_0 \\ \alpha_0 & \text{for } v, u \geqslant s_0\,. \end{cases}$$

π extends π_0 through repetition of (extensions of) the first row and first column. □

The diagonal restrictions. For the diagonal restrictions Ω_{ii} the size of a minimal realization may depend on the existence of fixed points under T other than δ_i. We show that the minimal size is equal to $|q_{ii}|$ if there are no such fixed points, and equal to the least even number greater than or equal to $|q_{ii}|$ otherwise. Put

$$d_i := \begin{cases} |q_{ii}| & \text{if } T\alpha \neq \alpha \text{ for all } \alpha \in q_{ii} \setminus \{\delta_i\} \\ 2\lceil \frac{1}{2}|q_{ii}|\rceil & \text{otherwise.} \end{cases}$$

Lemma 7.30. *If $q_{ii} \neq \{\delta_i\}$, then there are realizations $\pi\colon s \times s \to q_{ii}$ of Ω_{ii} exactly for all $s \geqslant d_i$.*

Proof. First we argue that $s \geqslant d_i$ is necessary. Trivially $s \geqslant |q_{ii}|$ is necessary, since each $\alpha \in q_{ii}$ has to occur at least once in every row and column. $d_i > |q_{ii}|$ if and only if $|q_{ii}|$ is odd and there is some $\alpha \neq \delta_i$ such that $T\alpha = \alpha$. In this case $d_i = |q_{ii}|+1$. But then this α has to occur an even number of times under π, whence either s has to be even (and therefore $s \geqslant |q_{ii}| + 1 = d_i$ in this case), or, if s is odd, α occurs at least twice in at least one row. This row still has to realize all other elements of q_{ii}, and it follows that $s \geqslant |q_{ii}| + 1 = d_i$ in that case as well.

Now for the existence of realizations as claimed. First consider realizations over $d_i \times d_i$.

If $|q_{ii}|$ is even or if $T\alpha \neq \alpha$ for all $\alpha \in q_{ii} \setminus \{\delta_i\}$, consider $(\Omega_{ii}; \nu)$, with ν identically put to 1, as a weighted game tableau. $(\Omega_{ii}; \nu)$ satisfies the requirements of Lemma 7.26 so that we obtain a realization on $d_i \times d_i$ since $d_i = |q_{ii}| = \sum_{\alpha \in q_{ii}} \nu(\alpha)$.

Otherwise $|q_{ii}|$ is odd, there is some $\alpha \neq \delta_i$ with $T\alpha = \alpha$, and $d_i = |q_{ii}|+1$. Let α_0 be the least $\alpha \in q_{ii} \setminus \{\delta_i\}$ that is fixed by T. Put $\nu(\alpha_0) = 2$ and $\nu(\alpha) = 1$

for all $\alpha \neq \alpha_0$. Again, $\sum_{\alpha \in q_{ii}} \nu(\alpha) = d_i$ is even and Lemma 7.26 applies to give a realization on $d_i \times d_i$.

From these minimal realizations on $d_i \times d_i$ one may again obtain realizations over $s \geqslant d_i$ simply by extensions that essentially repeat one row and one column. Let $\pi_0 : d_i \times d_i \to q_{ii}$ be the minimal realization, let α_0 be the minimal element of $q_{ii} \setminus \{\delta_i\}$. Then the following is a realization over $s \geqslant d_i$:

$$\pi(u,v) := \begin{cases} \pi_0(u,v) & \text{for } u,v < d_i \\ \pi_0(d_i-1,v) & \text{for } u \geqslant d_i, v < d_i - 1 \\ \pi_0(u,d_i-1) & \text{for } v \geqslant d_i, u < d_i - 1 \\ \delta_i & \text{for } u = v \geqslant d_i \\ \alpha_0 & \text{for } u > v \geqslant d_i - 1 \\ T\alpha_0 & \text{for } v > u \geqslant d_i - 1\,. \end{cases}$$

Figure 7.7 illustrates this extension of the domain by one row and one column to $s = d_i + 1$. □

Fig. 7.7

α	β	γ	$\cdots$	α_0	δ_i
α	β	γ	$\cdots$	δ_i	$T\alpha_0$
				$\vdots$	$\vdots$
				$T\gamma$	$T\gamma$
				$T\beta$	$T\beta$
				$T\alpha$	$T\alpha$

It remains to determine the size of a minimal realization for the entire game tableau Ω. Put

$$n_{ij} := \begin{cases} |q_{ij}| & \text{for } i \neq j \\ |q_{ii}| & \text{if } i = j \text{ and } T\alpha \neq \alpha \text{ for all } \alpha \in q_{ii} \setminus \{\delta_i\} \\ 2\lceil \frac{1}{2}|q_{ii}| \rceil & \text{if } i = j \text{ and } T\alpha = \alpha \text{ for some } \alpha \in q_{ii} \setminus \{\delta_i\} \end{cases} \tag{7.3}$$

so that according to the last lemmas n_{ij} is the minimal number such that Ω_{ij} admits a realization over $n_{ij} \times n_{ij}$. Note that for any game tableaux $n_{ij} = n_{ji}$ because T acts as a fixed-point free bijection between q_{ij} and q_{ji}. Assume that the q_{ij} are all non-empty. Then

- each Ω_{ij} for $i \neq j$ admits a realization over $s \times t$ for all $s,t \geqslant n_{ij}$,
- each non-trivial Ω_{ii} admits a realization on $s \times s$ for all $s \geqslant n_{ii}$, and

– each trivial $\mathfrak{Q}_{ii}$, $|q_{ii}| = \{\delta_i\}$, admits only the trivial singleton realization.

Optimal values for realizations that fit together in the sense of the condition in Proposition 7.20 are therefore given by:

$$\begin{array}{rcl} n_i & := & \max\{n_{ij} \mid 1 \leqslant j \leqslant l\} \\ n & := & \sum_i n_i \,. \end{array} \tag{7.4}$$

Notice that (ii) of Proposition 7.28 implies that $n_i = 1$ whenever $|q_{ii}| = 1$. If $\mathfrak{Q}$ therefore satisfies all conditions of Proposition 7.28 we do get a realization of size n.

Proposition 7.31. *Let $\mathfrak{Q}$ be a game tableau that satisfies the conditions of Proposition 7.28. Then either $|q_{ii}| = 1$ for all i and the only realization of $\mathfrak{Q}$ is over $n = \sum n_i = |\Delta|$; or there is at least one $q_{ii} \neq \{\delta_i\}$ and in this case $\mathfrak{Q}$ has realizations exactly over all $s \geqslant n$. (The n_i and n are as determined by equations 7.3 and 7.4.)*

The explicit constructions of realizations for the individual $\mathfrak{Q}_{ij}$ presented above and the general procedure for the composition of these according to Proposition 7.20 yield a PTIME algorithm as required for Proposition 7.28. This finishes the proof of Theorem 7.14.

It might be interesting to find a simple bound on the size n of a minimal realization of $\mathfrak{Q}$ also in terms of $|Q| = \sum_{1\leqslant i,j\leqslant l} |q_{ij}|$. Recall that $|Q|$ is the size of the L^2-invariant of the desired structure, or — in more model theoretic terms — the number of distinct L^2-types of pairs that the desired structure has to realize. We claim that actually $|Q| + 1$ is such a bound. In particular this is a *linear bound*, whereas for $k \geqslant 3$ we know that there cannot even be a sub-exponential bound on the size of minimal realizations of L^k-invariants by Example 3.23.

The following proposition gives a somewhat tighter bound in terms of both the number of 2-types and the number of 1-types that are to be realized.

Proposition 7.32. *For all $\mathfrak{A} \in \mathrm{fin}[\tau]$ there is some $\mathfrak{B} \in \mathrm{fin}[\tau]$ such that*

$$\mathfrak{B} \equiv^{L^2} \mathfrak{A} \quad \textit{and} \quad |B| \leqslant |\mathrm{Tp}^{L^2}(\mathfrak{A};2)| + 1 - \left(|\mathrm{Tp}^{L^2}(\mathfrak{A};1)| - 1\right)^2.$$

In particular $|B| \leqslant |\mathrm{Tp}^{L^2}(\mathfrak{A};2)| + 1 = |I_{L^2}(\mathfrak{A})| + 1$.

Proof. Let $\mathfrak{Q} = \left(Q, \leqslant, E, T, \Delta\right)$ be the game tableau associated with $I_{L^2}(\mathfrak{A})$. The desired $\mathfrak{B}$ is obtained from a realization of $\mathfrak{Q}$ of minimal size. Let $q_1, \ldots, q_l$, the q_{ij}, n_{ij}, n_i and n be as described in Proposition 7.28 and equations 7.3 and 7.4.

As $|Q| = |\mathrm{Tp}^{L^2}(\mathfrak{A};2)|$ and $|\mathrm{Tp}^{L^2}(\mathfrak{A};1)| = |Q/E| = l$ it suffices to show

$$n \leqslant |Q| + 1 - (l-1)^2,$$

since n is the minimal size of a realization of $\mathfrak{Q}$ as determined above. We first observe that $n_i = \max\{n_{ij} | 1 \leqslant j \leqslant l\} \leqslant \sum_{1 \leqslant j \leqslant l} n_{ij} - (l-1)$ because all n_{ij} are positive. Recall that $n_{ij} \leqslant |q_{ij}| + 1$ and $n_{ij} = |q_{ij}|$ at least for all $i \neq j$. Therefore

$$\begin{aligned} n = \textstyle\sum_i n_i &\leqslant \textstyle\sum_{i,j} n_{ij} - l(l-1) \\ &\leqslant \textstyle\sum_{i \neq j} |q_{ij}| + \sum_i \big(|q_{ii}| + 1\big) - l(l-1) \\ &\leqslant \textstyle\sum_{i,j} |q_{ij}| + l - l(l-1) = |Q| + 1 - (l-1)^2, \end{aligned}$$

as desired. □

The given bound is essentially optimal among bounds that are independent of the vocabulary. This is demonstrated in the following example.

Example 7.33. Let $k \geqslant 2$ and let τ_k consist of k binary relation symbols R_0, $\ldots$, R_{k-1}. Let $\mathfrak{A}_k$ consist of $2k-2$ points arranged in a cycle and with R_i interpreted by the set of pairs at distance i. The following sentence axiomatizes the complete L^2-theory of $\mathfrak{A}_k$.

$$\begin{aligned} \varphi_k \quad = \quad & \forall x \forall y \Big(\bigvee_i R_i xy \wedge \bigwedge_{i \neq j} \neg (R_i xy \wedge R_j xy) \Big) \\ & \wedge \forall x \forall y \Big(R_0 xy \leftrightarrow x = y \wedge \bigwedge_i (R_i xy \leftrightarrow R_i yx) \Big) \\ & \wedge \forall x \bigwedge_i \exists y R_i xy\,. \end{aligned}$$

Models of φ_k exactly correspond to realizations of the game tableau $\big(\mathfrak{Q}_k, \leqslant, E, T, \Delta\big)$ where $Q_k = k$, $E = k \times k$, $T = \mathrm{id}_k$, $\Delta = \{0\}$. In fact, if $\mathfrak{B}$ is a τ_k-structure over universe n, then $\mathfrak{B} \models \varphi_k$ if and only if

$$\begin{aligned} \pi\colon n \times n &\longrightarrow Q_k \\ (b, b') &\longmapsto i \text{ if } (b, b') \in R_i^{\mathfrak{B}} \end{aligned}$$

is a realization of $\mathfrak{Q}_k$. By Lemma 7.30 φ_k has models exactly in sizes greater than or equal to $n = 2\lceil \frac{1}{2} k \rceil$. Clearly $|\mathrm{Tp}^{L^2}(\mathfrak{A}_k; 2)| = k$ and $|\mathrm{Tp}^{L^2}(\mathfrak{A}_k; 1)| = l = 1$, so that $|Q| + 1 - (l-1)^2 = k + 1 = n$ for all odd k.

For situations with $l > 1$ one obtains similar examples by considering structures

$$\mathfrak{A} = \big(\mathfrak{A}_{k_1} \dot\cup \ldots \dot\cup \mathfrak{A}_{k_l}, P_1, \ldots, P_l\big)$$

with extra unary P_i to encode the partition into the $\mathfrak{A}_{k_i}$. Here $|\mathrm{Tp}^{L^2}(\mathfrak{A}; 2)| = \sum k_i + l(l-1)$, $|\mathrm{Tp}^{L^2}(\mathfrak{A}; 1)| = l$ so that $|Q| + 1 - (l-1)^2 = \sum k_i + l$. This bound is shown to be exact as above if all k_i are odd.

Another, more simple corollary to our findings about realizations concerns the spectrum of complete $L^2_{\infty\omega}$-theories.

Corollary 7.34. *Any complete $L^2_{\infty\omega}$-theory (in a finite relational vocabulary) that has any finite models, either has exactly one finite model up to isomorphism, or has models exactly in all cardinalities above some finite threshold value n.*

Sketch of Proof. Completeness of the theory together with existence of at least one finite model implies that all models (in fact finite and infinite ones) have the same value for their L^2-invariant. We know from Theorem 7.12 that each finite model of the theory is obtained from a realization of the underlying game tableau and vice versa. For realizations the corresponding spectrum property is expressed in Proposition 7.31 above. In the case of at least one non-trivial q_{ii} the above constructions for the extensions of realizations can easily be extended to yield models in arbitrary infinite cardinalities as well. □

7.3.2 On the Special Nature of Two Variables

Combinatorially, and with respect to the solution of the inversion problem presented here, the two-variable invariants and their induced tableaux are special. The trivialization of the accessibility relations E_1 and E_2 and modularity of the solutions as discussed in Section 7.1.1 are peculiar to the two-dimensional case in this sharp form. Basically the easy decomposition can be attributed to the fact that the two-variable type of a pair is fully determined by the individual two-variable types of its components together with the atomic type of the pair — a property that technically is reflected in Lemma 7.16. Combinatorially more sophisticated techniques may be required to approach the three-variable case. We have no well-founded conjecture at this stage whether indeed the k-variable case can be settled positively for any $k \geqslant 3$. In view of the general theorems above these canonization and inversion problems with and without counting remain challenging open problems.

The most important aspect with respect to classical logical concerns, in which two variables are very special, is decidability.

Theorem 7.35 (Mortimer). *$L^2_{\omega\omega}$ has the finite model property, i.e. any satisfiable sentence of $L^2_{\omega\omega}$ in a relational vocabulary has a finite model. Consequently, the* satisfiability problem *for $L^2_{\omega\omega}$ is decidable.*

Decidability of the satisfiability problem for $L^2_{\omega\omega}$ was earlier announced by Scott [Sco62], but the argument Scott gave was based on the erroneous assumption that the Gödel case with equality is decidable. So [Sco62] proves the claim only for $L^2_{\omega\omega}$ without equality. A version of Mortimer's proof [Mor75] can be found in [EF95]. There is also a new proof (with better complexity bounds) by Grädel, Kolaitis and Vardi [GKV96, BGG96].

Let us consider in a brief sketch what Mortimer's result implies about the inversion problem for I_{L^2}. To this end one may transform the information in a given candidate L^2-invariant

$$\mathfrak{I} = (Q, \leqslant, E, T; \Theta)$$

into an $L^2_{\omega\omega}$-sentence as follows (compare also Example 7.33). Introduce new predicates R_α for $\alpha \in Q$. Then the following sentence represents the full information in $\mathfrak{I}$. It may in fact be considered as an axiomatization of those realizations of the underlying game tableau that also respect atomic types as prescribed by Θ, whence it exactly axiomatizes the canonical R_α-expansions of all τ-structures $\mathfrak{A}$ with $I_{L^2}(\mathfrak{A}) = \mathfrak{I}$.

$$\begin{aligned}
&\forall x \forall y \bigvee_{\alpha} R_\alpha xy \ \wedge\ \forall x \forall y \bigwedge_{\alpha \neq \alpha'} \neg (R_\alpha xy \wedge R_{\alpha'} xy) \\
\wedge\ &\bigwedge_{\alpha} \forall x \forall y (R_\alpha xy \leftrightarrow R_{T\alpha} yx) \ \wedge \bigwedge_{\alpha \in \Delta, (\alpha, \alpha') \in E} \forall x (R_\alpha xx \leftrightarrow \exists y R_{\alpha'} xy) \\
\wedge\ &\bigwedge_{\alpha} \forall x \forall y \big(R_\alpha xy \rightarrow \Theta(\alpha)(x, y)\big).
\end{aligned}$$

This sentence may be regarded as a variant of a Scott sentence, with the crucial difference that it is over an extended vocabulary but requires only quantifier rank 2 and is of quadratic length in the size of the given invariant.

From the proof in [Mor75] one can infer that owing to the special format of this $L^2_{\omega\omega}$-sentence the size of its minimal models is bounded by a polynomial in the size of the given invariant.

The inversion problem thus reduces to the satisfiability problem for these associated compressed Scott sentences. This reduction induces an exponential time decision procedure for image(I_{L^2}) and a corresponding solution through exhaustive search to the inversion problem in exponential time.

Quite recently it has been shown in [GOR96b] that also the satisfiability problem for $C^2_{\omega\omega}$ is decidable (although $C^2_{\omega\omega}$ does not have the finite model property, compare Example 1.19).

Theorem 7.36 (Grädel, Otto, Rosen). *The satisfiability problem for $C^2_{\omega\omega}$ is decidable.*

It is remarkable on the other hand that even decidability of image(I_{L^k}) for $k > 2$ is an open problem. The corresponding problem for the I_{C^k} is trivial, since the size of candidate structures is easily determined from the proposed invariant. It should be stressed that also for the I_{L^k} there is no obvious connection between the decidability of the set of all invariants of actual structures and PTIME invertibility of I_{L^k} in the sense of Definition 6.9. In fact the size of prospective realizations may always in this context be thought of as a given parameter, in which case decidability becomes trivial through exhaustive search (compare the remarks following Definition 6.9).