Enhancing fixed point logic with cardinality quantifiers

Lauri Hella
Dept. of Mathematics
P.O. Box 4 (Yliopistonkatu 5)
00014 Univ. of Helsinki
Finland
hella@cc.helsinki.fi

Henrik Imhof
Dept. of Computer Science
Univ. of Wales, Swansea
Singleton Park
Swansea SA2 8PP
United Kingdom
H.Imhof@swan.ac.uk

Abstract

Let Q_{IFP} be any quantifier such that $FO(Q_{IFP})$, first order logic enhanced with Q_{IFP} and its vectorizations, equals inductive fixed point logic, *IFP* in expressive power. It is known that for certain quantifiers Q, the equivalence $FO(Q_{IFP}) \equiv IFP$ is no longer true if Q is added on both sides [12, 13]. Rather, we have $FO(Q_{IFP}, Q) < IFP(Q)$ in such cases. We extend these results to a great variety of quantifiers, namely all unbounded simple cardinality quantifiers.

Our argument also applies to partial fixed point logic, *PFP*. In order to establish an analogous result for least fixed point logic, *LFP*, we exhibit a general method to pass from arbitrary quantifiers to monotone quantifiers.

Our proof shows that the tree isomorphism problem is not definable in $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$, infinitary logic extended with all monadic quantifiers and their vectorizations, where a finite bound is imposed to the number of variables as well as to the number of nested quantifiers in $\mathbf{Q_1}$. This strengthens a result of Etessami and Immerman [5] by which tree isomorphism is not definable in $TC + COUNTING$.

1 Introduction

We study two well known mechanisms of increasing the expressive power of logics: adding additional quantifiers as opposed to the adjunction of recursive mechanisms (realized by fixed point operators). The following is known about the interaction of these two concepts. For any two logics, $\mathcal{L}_1, \mathcal{L}_2$, we write $\mathcal{L}_1 \leq \mathcal{L}_2$ to indicate that each $\varphi \in \mathcal{L}_1$ is equivalent to

a $\psi \in \mathcal{L}_2$. Derived from this, $\mathcal{L}_1 \equiv \mathcal{L}_2$ means that both $\mathcal{L}_1 \leq \mathcal{L}_2$ and $\mathcal{L}_2 \leq \mathcal{L}_1$ hold.
1. There is a class Q_{IFP} of structures such that $FO(Q_{IFP}) \equiv IFP$. Here we regard classes of structures as quantifiers. $FO(Q_{IFP})$ and IFP denote the extension of first order logic by Q_{IFP} and its vectorizations, and inductive fixed point logic, respectively. Such classes Q_{IFP} can be found in [2, 6, 10, 13].
2. By adjunction of further quantifiers one can make this equivalence break down. More precisely, one can find a class Q with $FO(Q_{IFP}, Q) < IFP(Q)$. This is possible even for linearly ordered structures [13]. For these structures, IFP, and hence $FO(Q_{IFP})$ captures $PTIME$. Moreover, $IFP(Q)$ captures $PTIME^Q$, the class of queries that are computable in polynomial time with access to the oracle Q (i.e. a string encoding of Q). However, for Q as above, $FO(Q_{IFP}, Q)$ does not capture all of $PTIME^Q$. This illustrates the important role of fixed point logic in descriptive complexity theory, and shows that it cannot be replaced by quantifier extensions of first order logic in some cases.

The class Q given in [13] for the second statement was obtained by a diagonal argument. Clearly, Q is not in $PTIME$. However, if the assumption on the structures of being ordered is dropped, the phenomenon $FO(Q_{IFP}, Q) < IFP(Q)$ occurs already for quantifiers of very low complexity, for instance for the even cardinality quantifier [12].

In this paper, we show that this phenomenon can be regarded as the usual case, in that it holds for a great variety of quantifiers Q. Indeed, we have $FO(Q_{IFP}, Q) < IFP(Q)$ for any unbounded simple cardinality quantifier (the definition is given below), and, in a probabilistic sense, almost all simple cardinality quantifiers are unbounded. The same holds if IFP is replaced with partial fixed point logic, PFP. In order to transfer the results to least fixed point logic, LFP, we pass from each Q to a monotone version Q^{mon}. As quantifiers, Q and Q^{mon} can be regarded as equivalent, since they will be first order definable from each other. If, as above, Q is an unbounded simple cardinality quantifier, we get $FO(Q_{LFP}, Q^{mon}) < LFP(Q^{mon})$.

2 Basic Definitions

Throughout this paper, we consider finite structures, only. Q will always denote a class of finite structures (that is closed under isomorphisms) over a finite relational vocabulary $\sigma = \{R_1, \ldots, R_s\}$. The arity of each relation R_i will be denoted r_i. Regarding Q as a quantifier allows to extend logics in a natural way [18]. We give a definition that adds to a logic $\mathcal{L}$ not only the capability of defining Q itself (as it was done by

Lindström in [18]), but also a uniform sequence of quantifiers derived from Q by means of vectorization. This approach has proved useful in the context of finite model theory, where quantifiers provide a mechanism of creating logics that capture certain complexity classes.

Definition 1: Let $\mathcal{L}$ be a logic the syntax of which is given by a collection of certain formation rules (involving first order variables $x_1, x_2, \ldots$). Then, for each vocabulary τ, $\mathcal{L}(Q)[\tau]$ is the smallest set of formulas that is closed under the formation rules for $\mathcal{L}$ enhanced, for each $r \geq 1$, with the following clause:

If $\overline{x}_i$ is an rr_i–tuple of variables $(1 \leq i \leq s)$ and $\Psi = \psi_1(\overline{x}_1), \ldots, \psi_s(\overline{x}_s)$ is a tuple of formulas of $\mathcal{L}(Q)[\tau]$, then

$$Q\overline{x}_1, \ldots, \overline{x}_s; \psi_1, \ldots, \psi_s$$

is again a formula in $\mathcal{L}(Q)[\tau]$. A τ–structure $\mathcal{A}$ satisfies this formula if and only if the σ–structure

$$\mathcal{A}_\Psi = (A^r; \psi_1^{\mathcal{A}}, \ldots, \psi_s^{\mathcal{A}})$$

belongs to Q. Here, $\psi_i^{\mathcal{A}}$ denotes the relation $\{\overline{a} \in A^{rr_i} \mid \mathcal{A} \models \psi_i[\overline{a}]\}$ for $1 \leq i \leq s$. This set is viewed as a set of r_i–tuples over A^r rather than rr_i–tuples over A. Thus $\mathcal{A}_{\overline{\psi}}$ is indeed a σ–structure. Note that in the above formulas ψ_i there can occur further free variables. They play the role of parameters. Extending the usual definition, a variable v is said to be free in $Q\overline{x}_1, \ldots, \overline{x}_s; \psi_1, \ldots, \psi_s$ if for some i it is free in ψ_i but is not contained in $\overline{x}_i$. □

In the sense of Lindström [18], we have added to $\mathcal{L}$ not only the class Q itself but also the following classes Q^r, which are derived from Q by vectorization:

Definition 2: Given $r \geq 1$, relation symbols S_i $(1 \leq i \leq s)$ of arity rr_i, let

$Q^r := \{(A; S_1, \ldots, S_s) \mid (A^r; S_1, \ldots, S_s) \in Q\}$,

where, on the right side, S_i is viewed as an r_i–ary relation on A^r. □

As an example, let $\sigma = \{E, S, T\}$ with unary S, T, and binary E, and take for Q those directed graphs $\mathcal{B}$ that have two nodes u, v, contained in the relations S, and T respectively, such that there is an E–path from u to v. Then, Q can be defined in *transitive closure logic*, TC, namely by the formula

$$\varphi = \exists uv(Su \wedge Tv \wedge [TC_{x,y}Exy]uv).$$

(For a definition and basic properties of TC, see [4], Section 7.6.) On the other hand, each formula $[TC_{\overline{x},\overline{y}}\chi]\overline{uv}$ of transitive closure logic is equivalent to the $FO(Q)$–formula

$$Q\overline{xy}, \overline{z}, \overline{w};\ \chi, \overline{z} = \overline{u}, \overline{w} = \overline{v}.$$

Here, vectorization was used to assert connectivity of a graph the vertices of which are tuples rather than single elements.

Enriching first order logic, FO, by one or several classes Q according to Definition 1 leads, in a sense, to a minimal logic making the classes Q^r expressible. To make this statement more precise, we consider regular logics $\mathcal{L}$ (cf. [3]), which are closed under first order operations and under substitutions of the form $R/\lambda\overline{x}\varphi(\overline{x})$ (with $\overline{x}$ and R of the same arity). That is, in any $\mathcal{L}$–formula involving a relation symbol R, one can replace R by a relation which is given by any other $\mathcal{L}$–formula φ. In general, regularity also includes closure under relativization but we will not use this property here.

Lemma 3 (Minimality Lemma): Let $\mathcal{L}$ be regular with $Q^r \in \mathcal{L}$ for all $r \geq 1$. Then $FO(Q) \leq \mathcal{L}$, that is, every $FO(Q)$–formula is equivalent to an $\mathcal{L}$–formula. Also, if several quantifiers (and their vectorizations) are adjoined to FO, this leads to a minimal logic in this sense.

Proof: We restrict to the case where one quantifier is added to FO. The proof is by induction on $FO(Q)$. In the Q-step, if Q occurs with arity r in a formula $Q\overline{x}_1, \ldots, \overline{x}_s; \psi_1, \ldots, \psi_s$, take an $\mathcal{L}$–formula defining Q^r and substitute the relations $S_1, \ldots, S_s$ (see Definition 2) by their 'definitions' $\psi_1, \ldots, \psi_s$, more precisely by $\mathcal{L}$–translations of these formulas. $\square$

It is straightforward to prove that any logic of the form $FO(Q)$ is *closed under vectorization*, i.e., if a class Q' is definable in $FO(Q)$, then so are the classes Q'^r, $r \geq 1$.

Next, we recall the **definition of fixed point logic:** *Inductive fixed point logic*, *IFP*, allows for second order variables $X, Y, \ldots$, which may only occur bounded by fixed point operators. The meaning of $[IFP_{\overline{x},X}\,\theta(\overline{x}, X)]\overline{t}$ is $\overline{t} \in X_\infty$, where we set $X_0 := \emptyset$, $X_{j+1} := X_j \cup \{\overline{x} \mid \theta(\overline{x}, X_j)\}$, $X_\infty := \bigcup_{j\geq 0} X_j$.

Partial fixed point logic, *PFP*, is defined similarly, setting however $X_{j+1} := \{\overline{x} \mid \theta(\overline{x}, X_j)\}$, and letting $X_\infty := X_j$ for the smallest j such that $X_j = X_{j+1}$ if such a j exists, and $X_\infty := \emptyset$ otherwise.

Least fixed point logic, *LFP*, is obtained by restricting the application of fixed point operators to the case where $\theta(\overline{x}, X)$ is positive in X.

It is easy to see that $LFP \leq IFP \leq PFP$. Gurevich and Shelah [7] proved that $IFP \equiv LFP$. In the presence of linear order, *IFP* captures *PTIME* ([14, 19]) and *PFP* captures *PSPACE* ([19, 1]).

We refer to Chapters 6 and 7 of [4] for a thorough treatment of these fixed point logics and their role in descriptive complexity theory.

Fact 4 [2, 6, 10, 13]: There is a class Q_{IFP} such that $FO(Q_{IFP}) \equiv IFP$ ($\equiv LFP$).[1] Similarly, there is a class Q_{PFP} such that $FO(Q_{PFP}) \equiv PFP$.

It is known that these equivalences do not survive the extensions with further quantifiers. That is, $FO(Q_{IFP}, Q) \equiv IFP(Q)$ fails for certain classes Q [13, 12]. By Lemma 3, the logic on the left side of this equation is minimal (and does, in particular, not depend on the choice of Q_{IFP}), whereas on the right side the interplay of fixed point operators with quantifiers may result in an increase of expressive power. In this paper, we show that for cardinality quantifiers this is even the usual case, that is, it occurs with probability one.

To make the problem more transparent, we characterize $FO(Q_{IFP}, Q)$ as a fragment of $IFP(Q)$, thereby exhibiting the role of second order variables in fixed point logics.

Remark 5 [13]: Let Q_{IFP} be a class of structures such that $IFP \equiv FO(Q_{IFP})$ holds. Then for any class Q, we have $FO(Q_{IFP}, Q) \equiv IFP(Q)^*$, where $IFP(Q)^*$ denotes the fragment of $IFP(Q)$ which is obtained by allowing the formation of $Q\overline{x}_1, \ldots, \overline{x}_s; \psi_1, \ldots, \psi_s$ only in case $\psi_1, \ldots, \psi_s$ have no *free* occurences of *second order* variables. Analogous results hold for PFP and LFP (in the latter case one has to restrict to a monotone Q, cf. Section 5).

3 Cardinality quantifiers and infinitary languages

If the signature σ of a quantifier Q consists of monadic relation symbols, only, then we refer to Q as a *cardinality quantifier*. This is because the Q–membership of a σ–structure $\mathcal{B}$ depends only on the cardinality of the Boolean combinations of these relations (respectively of $|B|$, if $\sigma = \emptyset$). We write $\mathbf{Q_1}$ for the collection of all cardinality quantifiers.

Given $Q \in \mathbf{Q_1}$, we are going to embed the logic $FO(Q_{IFP}, Q)$ into an infinitary language and define a game which respects the nesting of Q. This will enable us to establish queries that are not definable in this infinitary logic, but can be expressed in $IFP(Q)$, where bounded nesting

[1] For example, Q_{IFP} can be taken to be the class of all structures $(A; E, U, S, T)$, $E \subseteq A^2$, $U, S, T \subseteq A$, such that $S \times T$ is contained in the alternating transitive closure $ATC(E, U)$ (see, e.g., [4], Section 8.4).

can be overcome by storing information in the second order variable. Remember that, by Remark 5, it is exactly the possibility to make this information available to the quantifier Q which separates the two logics.

Let $\mathcal{L}_{\infty\omega}$ denote the extension of first order logic that allows for conjunctions and disjunctions over arbitrary sets of formulas. As each finite structure can be characterized up to isomorphism by a single first order formula, $\mathcal{L}_{\infty\omega}$ can express all queries on finite structures, and hence is of no interest. However, the restriction $\mathcal{L}^{\omega}_{\infty\omega}$, consisting of all formulas that involve only a finite number of distinct variables, has interesting model theoretic properties, and is still strong enough to simulate fixed point operations. Indeed, Kolaitis and Vardi [16] proved that $PFP \leq \mathcal{L}^{\omega}_{\infty\omega}$. In accordance with Definition 1, $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})$ denotes the extension of this logic with all cardinality quantifiers and their vectorized versions. However, it makes no difference whether we allow vectorization or not, as we will see below. To adapt this infinitary logic to our needs, we impose a restriction on the nesting of quantifiers.

Definition 6:

(i) Let, for $l \geq 0$, $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^l$ denote the collection of $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})$–formulas in which there occur at most l nested applications of (possibly vectorized) quantifiers in $\mathbf{Q_1}$. Here, the nesting (or the quantifier rank) is defined by induction so that the nesting of a formula $Q\overline{x};\Phi$ (for $Q \in \mathbf{Q_1} \setminus \{\exists,\forall\}$) is $m+1$, where m denotes the maximum nesting in the tuple the formulas Φ, regardless of the arities of Q in its occurences.

(ii) Let $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^l$ be defined as $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^l$ apart from the fact that vectorization is no longer allowed (that is to say, in Definition 1, the arity r has to be 1).

(iii) Set $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega} := \bigcup_{l>0} \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^l$ and $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^{\omega} := \bigcup_{l>0} \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^l$.

Let $\mathbf{C}$ be the class of all *counting quantifiers* $Q_i := \{(A;R) \mid |R| \geq i\}$, $i \geq 1$. Replacing the class $\mathbf{Q_1}$ by $\mathbf{C}$ in the definition above, we obtain the corresponding bounded nesting versions $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C})^{\omega}$ and $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{\omega}$ of the infinitary counting logic $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C})$.

Proposition 7:

(i) For all cardinality quantifiers Q, we have $FO(Q_{IFP}, Q) \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$.

(ii) $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega} \equiv \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^{\omega} \equiv \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{\omega}$.

Proof: The inclusion (i) is obtained from the Minimality Lemma. Therefore, note first that $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$ is a regular logic. The crucial point for this observation is the closure under substitution of predicates. In fact, substituting a predicate R in an $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{l}$–formula with an $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{l'}$–formula results in an $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{l+l'}$–formula.

Furthermore, by definition, all Q^r $(r \geq 1)$ are $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$–definable, as well as all vectorizations of Q_{IFP} are (since $IFP \leq PFP \leq \mathcal{L}^{\omega}_{\infty\omega}$). Hence the Minimality Lemma applies.

Clearly $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{\omega} \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^{\omega} \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$. Thus, to see (ii), it suffices to show that all the vectorizations Q^r of quantifiers $Q \in \mathbf{Q_1}$ are $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{\omega}$–definable. To this end, we argue as was done in [15], Proposition 2.18. For each structure $\mathcal{B}$ over the vocabulary of Q^r there is a sentence $\psi_{\mathcal{B}}$ of $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{r}$ that describes $\mathcal{B}$ up to cardinalities of Boolean combinations of its relations. Since membership in Q^r depends only on these cardinalities, we conclude that the $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{r}$–sentence $\bigvee_{\mathcal{B}\in Q^r} \psi_{\mathcal{B}}$ defines Q^r. □

To establish non–definability results for $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^{\omega}$, and hence for each logic $FO(Q_{IFP}, Q)$ where Q is a cardinality quantifier, one can apply the bijective games defined in [8, 9]. We adapt the definition for the case of bounded nesting.

Definition 8: Let $\mathcal{A} \equiv_{k,l} \mathcal{B}$ denote the equivalence of two structures, $\mathcal{A}$ and $\mathcal{B}$, with respect to $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1^*})^{l}$–formulas that involve at most k distinct variables.

To prove this equivalence for two given structures we will use the k, l–pebble game defined as follows.

Definition (of the k, l–game) 9: The board consists of a copy of each of the structures $\mathcal{A}$ and $\mathcal{B}$, and there are k pairs $(u_1, v_1), \ldots, (u_k, v_k)$ of pebbles available which are off the board in the beginning. The u_i always have to be placed on $\mathcal{A}$, the v_i on $\mathcal{B}$. There are two players, spoiler and duplicator. In each round of the game, the spoiler may choose between a *pebble move* and a *bijective move*, but he is not allowed to make more than l bijective moves during the whole game.

In a pebble move, he may pick up a pebble from either structure and place it elsewhere on this structure, or take an unused pebble from the supply and put it on the appropriate structure. Then, the duplicator picks up the corresponding pebble and places it on the other structure.

In a bijective move, the duplicator must choose a bijection $f : A \to B$. Then the spoiler picks up a pebble u_i and places it on an element $a \in A$. The game continues with the pebbles u_i and

v_i placed on a and $f(a)$ respectively. If there are no bijections $A \to B$, i.e. if $|A| \neq |B|$, then the spoiler wins at this point.

In the case $|A| = |B|$, the spoiler wins if, at any stage, the mapping $\overline{a} \mapsto \overline{b}$, where $\overline{a}$ consists of the pebbled elements in $\mathcal{A}$ and $\overline{b}$ consists of the corresponding pebbled elements in $\mathcal{B}$, is not a partial isomorphism $\mathcal{A} \to \mathcal{B}$. Otherwise, the game goes on for infinitely many rounds, and the duplicator wins.

We write $\mathcal{A} \simeq_{k,l} \mathcal{B}$ if the duplicator has a winning strategy in the k, l–game on $\mathcal{A}$ and $\mathcal{B}$.

Theorem 10: Let $\mathcal{A}$ and $\mathcal{B}$ be two structures of the same vocabulary. $\mathcal{A} \equiv_{k,l} \mathcal{B}$ if and only if $\mathcal{A} \simeq_{k,l} \mathcal{B}$.

Proof: It is clear that the proof in [9], given for $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q}_1^*)$, the extension of $\mathcal{L}^{\omega}_{\infty\omega}$ with all cardinality quantifiers (without vectorizations), can be carried over to the case of bounded nesting. $\square$

We now turn our attention to so called *simple* cardinality quantifiers that have a vocabulary consisting of only one relation symbol. So, assume $\sigma = \{R\}$ for a monadic R, and let Q be a class of σ–structures. Observe that we can identify Q with an infinite sequence of strings, $w_1^Q, w_2^Q, \ldots$, over the alphabet $\{+, -\}$, defined as follows. w_n^Q is of length $n+1$, and in $w_n^Q = d_0 \ldots d_n$, the symbol d_m is a $+$ just in case the structure $\mathcal{A}$ of size n with $|R^{\mathcal{A}}| = m$ is in Q (for $0 \leq m \leq n$).

Definition 11: Let Q be a simple cardinality quantifier. For $s \geq 0$, we say that Q is *s–bounded*, if for all $n > 2s$, the string w_n^Q is of the form

$$w_n^Q = a_1 \ldots a_s \, v(Q,n)^{n+1-2s} \, b_1 \ldots b_s, \text{ for some } v(Q,n) \in \{+,-\}.$$

In that case, we call $v(Q,n)$ the *n–value*, and $p(Q,n,s) := a_1 \ldots a_s b_1 \ldots b_s$ the *n,s–pattern* of Q.

We call Q *bounded* if it is *s–bounded* for some $s \geq 0$, *unbounded* otherwise.

The notion of unboundedness was first defined in [15] for monotone simple cardinality quantifiers. In [11], it was proved that it is a key property for the definability of certain polyadic lifts of monotone simple cardinality quantifiers: the branching $Br(Q,Q)$ and the k–ary Ramseyfication $Ram^k(Q)$ of Q are definable in $FO(\mathbf{Q}_1^*)$ if and only if Q is bounded (see [11] for definitions of the lifts Br and Ram).

Note that a simple cardinality quantifier Q is unbounded if and only if for every s there are n and t such that

(1) $s \leq t \leq n - s$, and

(2) $d_t \neq d_{t+1}$, where $w_n^Q = d_0 \dots d_n$.

We say that Q is $(-+)$–unbounded ($(+-)$–unbounded) if the inequality $d_t \neq d_{t+1}$ in (2) can be replaced with the condition: $d_t = -$ and $d_{t+1} = +$ ($d_t = +$ and $d_{t+1} = -$, respectively). Clearly every unbounded quantifier Q is either $(-+)$–unbounded, or $(+-)$–unbounded (or both). Moreover, Q is $(-+)$–unbounded if and only if its complement $\neg Q$ is $(+-)$–unbounded.

A simple cardinality quantifier, Q, can be chosen at random by tossing a coin infinitely many often so as to determine the symbols in the infinite string $w_1^Q w_2^Q \dots$ For each $s \geq 0$, the probability of Q being s–bounded is zero because

$$\mathrm{Prob}[w_n^Q = a_1 \dots a_s v^{n+1-2s} b_1 \dots b_s \text{ for some } a_1 \dots a_s b_1 \dots b_s \text{ and } v] = (1/2)^{n-2s}.$$

So we get:

Theorem 12: Almost all simple cardinality quantifiers are unbounded. $\square$

Popular examples of unbounded quantifiers are the even cardinality quantifier, $Q_{EVEN} := \{(A; R) \mid R \text{ is of even cardinality}\}$, and the majority quantifier, $Q_{MAJ} := \{(A; R) \mid |R| \geq 1/2|A|\}$. But being bounded does not amount to being first order definable or being trivial. For instance, $Q_{EVENDOMAIN} := \{(A; R) \mid A \text{ is of even cardinality}\}$ is bounded. Similarly, any cardinality property of the *domain* can be represented this way. This includes nonrecursive properties.

Theorem 13: If Q is a bounded simple cardinality quantifier, then $FO(Q_{IFP}, Q) \equiv IFP(Q)$.

Proof: Choose an s such that Q is s–bounded. $P := \{+, -\}^{2s}$ is the set of all possibly occuring n, s–patterns. For each $p \in P$ and $r \geq 1$ there is a formula $\chi_{p,r} \in FO(Q)$ that holds in a structure of size n just in case $p(Q, n^r, s) = p$. Similarly, there is a formula $\chi_r^+ \in FO(Q)$ which says in cardinality n that $v(Q, n^r) = +$.

Let now φ be any $IFP(Q)$–formula. Consider its subformulas of the type $Q\overline{x}\psi(\overline{x})$, where ψ may contain parameters different from $\overline{x}$. For a given $p \in P$, this subformula is true in a structure $\mathcal{A}$ with $p(Q, |A|^r, s) = p$ and $v(Q, |A|^r) = +$ if and only if the number m of r-tuples over A that satisfy ψ is either in $\{s, \dots, n^r - s\}$ or in the subset of $\{0, \dots, s-1, n^r - s + 1, \dots, n^r\}$ that is given by the symbols $+$ in p. Hence there is a formula $\varphi^{+}_{p,r} \in \mathit{IFP}$ which is equivalent to φ in all structures $\mathcal{A}$ with $p(Q, |A|^r, s) = p$ and $v(Q, |A|^r) = +$. Let $\varphi^{-}_{p,r}$ be the analogous formula for the case $v(Q, |A|^r) = -$.

If we now replace $\varphi^{+}_{p,r}$ and $\varphi^{-}_{p,r}$ with their translations in $FO(Q_{\mathit{IFP}})$, we find that φ is equivalent to the $FO(Q_{\mathit{IFP}}, Q)$–formula

$$\bigvee_{p \in P} (\chi_{p,r} \wedge ((\chi^{+}_{r} \wedge \varphi^{+}_{p,r}) \vee ((\chi^{-}_{r} \wedge \varphi^{-}_{p,r}))). \quad \Box$$

Our main result, to be proved in the next section, is the converse of the previous theorem.

Main Theorem 14: Let Q be an unbounded simple cardinality quantifier. Then $FO(Q_{\mathit{IFP}}, Q) < \mathit{IFP}(Q)$.

The so called *random oracle hypothesis* says that two classes of queries are different if they are separated by almost all oracles. If one restricts to simple cardinality quantifiers then the above result together with Theorem 12 can be seen as a counterexample to a logical variant of that hypothesis.

4 Proof of the main result

This section is devoted to the proof of Theorem 14. We will first show that for all natural numbers k and l there are nonisomorphic trees $\mathcal{A}$ and $\mathcal{B}$ such that $\mathcal{A} \equiv_{k,l} \mathcal{B}$. By a *tree* we mean a directed graph $\mathcal{C} = (C; E^{\mathcal{C}})$ such that the edge relation $E^{\mathcal{C}}$ is cycle-free and every element of C, except one (the root of $\mathcal{C}$), has exactly one $E^{\mathcal{C}}$–predecessor. The *subtree generated by an element* $a \in C$ is $\mathcal{C}_a = (C_a; E^{\mathcal{C}_a})$, where $C_a = \{a\} \cup \{b \in C \mid \text{there is a directed } E^{\mathcal{C}}\text{–path from } a \text{ to } b\}$ and $E^{\mathcal{C}_a} = E^{\mathcal{C}} \cap (C_a)^2$.

The $\equiv_{k,l}$–equivalent trees will be obtained by iterating the following basic constructions:

Definition 15: Let k be a natural number, and let $\mathcal{A} = (A; E^{\mathcal{A}})$ and $\mathcal{B} = (B; E^{\mathcal{B}})$ be trees.
(i) For each i and j, $H^{i,j}(\mathcal{A}, \mathcal{B})$ is a tree which consists of a new root that is connected by an E–edge to the roots of i isomorphic copies of $\mathcal{A}$ and j isomorphic copies of $\mathcal{B}$, all mutually disjoint.
(ii) $F^k(\mathcal{A}, \mathcal{B})$ is a tree which is obtained by connecting a new root by an E–edge to the roots of one copy of $H^{k,k+2}(\mathcal{A}, \mathcal{B})$ and one copy of $H^{k+2,k}(\mathcal{A}, \mathcal{B})$.
(iii) $G^k(\mathcal{A}, \mathcal{B})$ is similar to $F^k(\mathcal{A}, \mathcal{B})$, except that instead of the copies of $H^{k,k+2}(\mathcal{A}, \mathcal{B})$ and $H^{k+2,k}(\mathcal{A}, \mathcal{B})$ it has two copies of $H^{k+1,k+1}(\mathcal{A}, \mathcal{B})$.

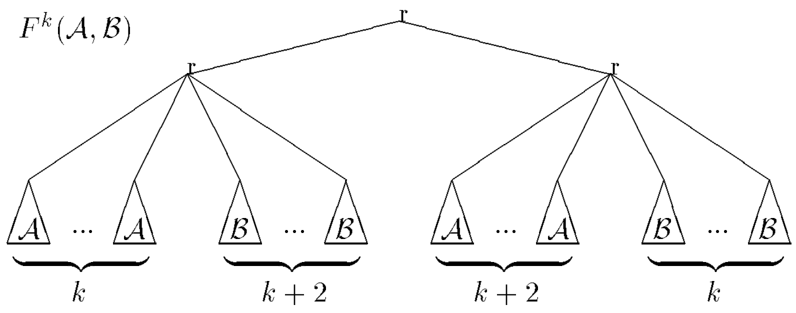

Note that both $F^k(\mathcal{A},\mathcal{B})$ and $G^k(\mathcal{A},\mathcal{B})$ contain $2k+2$ copies of $\mathcal{A}$ and the same number of copies of $\mathcal{B}$. Furthermore, it is easy to see that $F^k(\mathcal{A},\mathcal{B})$ and $G^k(\mathcal{A},\mathcal{B})$ are nonisomorphic if $\mathcal{A}$ and $\mathcal{B}$ are.

In the next two lemmas we will consider a slightly modified version of the k,l–game. The rules of this modified k,l–game are as given in Definition 9, except that in addition the spoiler always has to start with a bijective move. We write $\mathcal{C} \simeq^*_{k,l} \mathcal{D}$ if the duplicator has a winning strategy in the modified k,l–game on $\mathcal{C}$ and $\mathcal{D}$. Note that clearly $\mathcal{C} \simeq^*_{k,l+1} \mathcal{D}$ implies $\mathcal{C} \simeq_{k,l} \mathcal{D}$.

Lemma 16: For any trees $\mathcal{A}$ and $\mathcal{B}$, $F^k(\mathcal{A},\mathcal{B}) \simeq^*_{k,1} G^k(\mathcal{A},\mathcal{B})$.

Proof: For the sake of simplicity, we denote $F^k(\mathcal{A},\mathcal{B})$ by $\mathcal{C}$ and $G^k(\mathcal{A},\mathcal{B})$ by $\mathcal{D}$. Let a and a' be the roots of $\mathcal{C}$ and $\mathcal{D}$, respectively. Furthermore, let b and c be the roots of the copies of $H^{k,k+2}(\mathcal{A},\mathcal{B})$ and $H^{k+2,k}(\mathcal{A},\mathcal{B})$ in $\mathcal{C}$, and similarly, let b' and c' be the roots of the two copies of $H^{k+1,k+1}(\mathcal{A},\mathcal{B})$ in $\mathcal{D}$.

As observed above, $\mathcal{C}$ and $\mathcal{D}$ contain an equal number of copies of both $\mathcal{A}$ and $\mathcal{B}$: In $\mathcal{C}_b$ there are k copies $\mathcal{A}_1,\ldots,\mathcal{A}_k$ of $\mathcal{A}$ and $k+2$ copies $\mathcal{B}_1,\ldots,\mathcal{B}_{k+2}$ of $\mathcal{B}$, and in $\mathcal{C}_c$ there are $k+2$ copies $\mathcal{A}_{k+1},\ldots,\mathcal{A}_{2k+2}$ of $\mathcal{A}$ and k copies $\mathcal{B}_{k+3},\ldots,\mathcal{B}_{2k+2}$ of $\mathcal{B}$. Similarly, in $\mathcal{D}_{b'}$ there are $k+1$ copies $\mathcal{A}'_1,\ldots,\mathcal{A}'_{k+1}$ of $\mathcal{A}$ and $k+1$ copies $\mathcal{B}'_1,\ldots,\mathcal{B}'_{k+1}$ of $\mathcal{B}$, and in $\mathcal{D}_{c'}$ there are $k+1$ copies $\mathcal{A}'_{k+2},\ldots,\mathcal{A}'_{2k+2}$ of $\mathcal{A}$ and $k+1$ copies $\mathcal{B}'_{k+2},\ldots,\mathcal{B}'_{2k+2}$ of $\mathcal{B}$. Thus, the duplicator can start the modified k,l–game with a bijection f which maps a to a', b to b', c to c', and, for each $1 \le i \le 2k+2$, $\mathcal{A}_i$ and $\mathcal{B}_i$ isomorphically to $\mathcal{A}'_i$ and $\mathcal{B}'_i$, respectively.

Let $e \in C$ be the first move of the spoiler. If e is in neither of the subtrees $\mathcal{A}_{k+1}$ and $\mathcal{B}_{k+2}$, then the duplicator can play after the first round in such a way that a, b and c are always mapped to a', b' and c', respectively, and each copy of $\mathcal{A}$ in $\mathcal{C}_b$ ($\mathcal{C}_c$) is always mapped isomorphically to a copy of $\mathcal{A}$ in $\mathcal{D}_{b'}$ ($\mathcal{D}_{c'}$), and similarly for copies of $\mathcal{B}$. This is possible, because there are no more bijective moves, and whenever the spoiler picks up a pebble to play a pebble move, in each of the subtrees

$\mathcal{C}_b$, $\mathcal{C}_c$, $\mathcal{D}_{b'}$ and $\mathcal{D}_{c'}$ there are always at least one copy of $\mathcal{A}$ and one copy of $\mathcal{B}$ not containing any pebbles. On the other hand, if $e \in A_{k+1} \cup B_{k+2}$, then the duplicator has to map b to c' and c to b' in the continuation of the game, but otherwise he can play exactly as described above, since $\mathcal{D}_{b'}$ and $\mathcal{D}_{c'}$ are isomorphic. In both cases it is clear that the duplicator wins the game. □

Lemma 17: If $\mathcal{A}$ and $\mathcal{B}$ are trees such that $\mathcal{A} \simeq^*_{k,l} \mathcal{B}$, then $F^k(\mathcal{A}, \mathcal{B}) \simeq^*_{k,l+1} G^k(\mathcal{A}, \mathcal{B})$.

Proof: We use the same notation as in the preceding proof: $\mathcal{C} = F^k(\mathcal{A}, \mathcal{B})$, $\mathcal{D} = G^k(\mathcal{A}, \mathcal{B})$, a (a') is the root of $\mathcal{C}$ $(\mathcal{D})$, and b and c (b' and c') are the E–successors of a (a').

Assume that $\mathcal{A} \simeq^*_{k,l} \mathcal{B}$, and consider the modified $k, l+1$–game on $\mathcal{C}$ and $\mathcal{D}$. The duplicator can use exactly the same strategy as in the proof of Lemma 16, until the spoiler chooses to play a bijective move for the second time. At that point the duplicator can choose a bijection g such that $g(a) = a'$, $g(b) = b'$ and $g(c) = c'$ (or $g(b) = c'$ and $g(c) = b'$, depending on the first move of the spoiler as explained in the proof of Lemma 16), and which still maps all copies of $\mathcal{A}$ and $\mathcal{B}$ in $\mathcal{C}_b$ $(\mathcal{C}_c)$ that contain pebbles isomorphically to the corresponding copies of $\mathcal{A}$ and $\mathcal{B}$ in $\mathcal{D}_{b'}$ $(\mathcal{D}_{c'})$. Note that g has to map $\mathcal{C}_b$ to $\mathcal{D}_{b'}$ and $\mathcal{C}_c$ to $\mathcal{D}_{c'}$, whence there is one copy $\mathcal{A}_i$ of $\mathcal{A}$ in $\mathcal{C}$ that is mapped to a copy $\mathcal{B}'_j$ of $\mathcal{B}$ in $\mathcal{D}$, and vice versa. However, since $\mathcal{A}_i$ does not contain any pebbles, the duplicator can use the first bijection of his winning strategy in the modified k, l–game on $\mathcal{A}_i$ and $\mathcal{B}'_j$ in defining g on $\mathcal{A}_i$. Thus, for the rest of the game the duplicator can play in such a way that subtrees of $\mathcal{C}$ are mapped to subtrees of $\mathcal{D}$ either isomorphically, or according to a winning strategy in the modified k, l–game. It is clear that with this strategy the duplicator is guaranteed to win the modified $k, l+1$–game on $\mathcal{C}$ and $\mathcal{D}$. □

We define now for each $k \geq 1$ and $l \geq 0$ a pair $\mathcal{A}_{k,l}, \mathcal{B}_{k,l}$ of trees by induction on l as follows:

Definition 18:

(i) $\mathcal{A}_{k,0}$ is a tree which consists of a root a and a single leaf a': $A_{k,0} = \{a, a'\}$, $E^{\mathcal{A}_{k,0}} = \{(a, a')\}$. $\mathcal{B}_{k,0}$ is the trivial one-element tree: $B_{k,0} = \{b\}$, $E^{\mathcal{B}_{k,0}} = \emptyset$.

(ii) $\mathcal{A}_{k,l+1} = F^k(\mathcal{A}_{k,l}, \mathcal{B}_{k,l})$ and $\mathcal{B}_{k,l+1} = G^k(\mathcal{A}_{k,l}, \mathcal{B}_{k,l})$.

Using Lemmas 16 and 17 it is straightforward to prove by induction on $l \geq 1$ that $\mathcal{A}_{k,l} \simeq^*_{k,l} \mathcal{B}_{k,l}$. Hence, we have

Proposition 19: For all $k \geq 1$ and $l \geq 0$, $\mathcal{A}_{k,l+1} \equiv_{k,l} \mathcal{B}_{k,l+1}$. □

The next step in the proof of Theorem 14 is to show that there is an inductive property that separates the trees $\mathcal{A}_{k,l}$ and $\mathcal{B}_{k,l}$. In order to make this property definable in logics of the form $IFP(Q)$ we will adjust the cardinalities of $\mathcal{A}_{k,l}$ and $\mathcal{B}_{k,l}$ by adding sets of new elements (the same trick was also used in [11]):

Definition 20: Let p and q be natural numbers.

(i) $\mathcal{C}^{p,q} = (C; P)$ is a structure such that $|C| = p + q$ and $|P| = p$.

(ii) $\mathcal{A}^{p,q}_{k,l}$ is the $\{E, P\}$–structure which is obtained by taking the disjoint union of $\mathcal{C}^{p,q}$ and $\mathcal{A}_{k,l}$. Similarly, $\mathcal{B}^{p,q}_{k,l}$ is the disjoint union of $\mathcal{C}^{p,q}$ and $\mathcal{B}_{k,l}$.

It is easy to see that Proposition 19 remains valid for the extended structures $\mathcal{A}^{p,q}_{k,l+1}$ and $\mathcal{B}^{p,q}_{k,l+1}$. Indeed, the duplicator can extend his winning strategy in the k, l–game on $\mathcal{A}_{k,l+1}$ and $\mathcal{B}_{k,l+1}$ to a winning strategy on $\mathcal{A}^{p,q}_{k,l+1}$ and $\mathcal{B}^{p,q}_{k,l+1}$ just by using the identity function id_C on C.

Proposition 21: For all $p, q, k, l \geq 0$, $\mathcal{A}^{p,q}_{k,l+1} \equiv_{k,l} \mathcal{B}^{p,q}_{k,l+1}$. □

Let Q be an unbounded simple cardinality quantifier. We may assume without loss of generality that Q is $(-+)$–unbounded; otherwise its complement $\neg Q$ is $(-+)$–unbounded, and we can just replace Q by $\neg Q$ in what follows. Let $\theta(x, X)$ be the formula

$$\psi(x) \vee \exists y \chi(x, y, X),$$

where $\psi(x)$ and $\chi(x, y, X)$ are the formulas

$$\exists y(Exy \wedge \forall z(Exz \rightarrow z = y))$$

and

$$Exy \wedge \exists z(Eyz \wedge Xz) \wedge Qz(Pz \vee (Eyz \wedge Xz)),$$

respectively. Furthermore, for $\mathcal{D} = \mathcal{A}^{p,q}_{k,l}, \mathcal{B}^{p,q}_{k,l}$, let $X(\mathcal{D})_\infty$ be the inductive fixed point in $\mathcal{D}$ of the operator corresponding to the formula $\theta(x, X)$; that is, $X(\mathcal{D})_\infty = \bigcup_{j \geq 0} X(\mathcal{D})_j$, where $X(\mathcal{D})_0 = \emptyset$ and $X(\mathcal{D})_{j+1} = X(\mathcal{D})_j \cup \{a \in D \mid \mathcal{D} \models \theta[a, X(\mathcal{D})_j]\}$. Our aim is to show that for every k and l there are p and q such that the root of the tree $\mathcal{A}_{k,l}$ is in $X(\mathcal{A}^{p,q}_{k,l})_\infty$, but the root of the tree $\mathcal{B}_{k,l}$ is not in $X(\mathcal{B}^{p,q}_{k,l})_\infty$.

Let k and l be given, and let m be the cardinality of the trees $\mathcal{A}_{k,l}$ and $\mathcal{B}_{k,l}$. Since Q is $(-+)$–unbounded, we can choose n and t such that

(1) $k + 1 \leq t \leq n - m + k + 1$, and

(2) $d_t = -$ and $d_{t+1} = +$, where $w^Q_n = d_0 \ldots d_n$.

We set $p = t-(k+1)$ and $q = n-m-t+(k+1)$. Note that by condition (1) above, $p, q \geq 0$. Note also that the cardinality of the structures $\mathcal{A}^{p,q}_{k,l}$ and $\mathcal{B}^{p,q}_{k,l}$ is $m+p+q = n$.

Let $\mathcal{D}$ be either of the structures $\mathcal{A}^{p,q}_{k,l}, \mathcal{B}^{p,q}_{k,l}$, and consider the stages $X(\mathcal{D})_j$ of the induction that corresponds to the formula $\theta(x, X)$. We prove by induction on $j \geq 1$ that

$$X(\mathcal{D})_j = \{a \in D \mid \exists i < j : \mathcal{D}_a \text{ is isomorphic to } \mathcal{A}_{k,i}\}.$$

(i) From the definition of $\theta(x, X)$ we see that $a \in X(\mathcal{D})_1$ if and only if $\mathcal{D} \models \psi[a]$, i.e., if and only if a has exactly one E–successor. Clearly this holds if $\mathcal{D}_a$ is isomorphic to $\mathcal{A}_{k,0}$, but in all other cases a has either 0, 2 or $2k+2$ E–successors. Hence the claim holds in the case $j = 1$.

(ii) Assume that the claim holds for j. Suppose first that $\mathcal{D}_a$ is isomorphic to $\mathcal{A}_{k,i}$ for some $i < j+1$. If $i < j$, then, by induction hypothesis, $a \in X(\mathcal{D})_j \subseteq X(\mathcal{D})_{j+1}$. If $i = j$, then, by the definition of $\mathcal{A}_{k,i} = F^k(\mathcal{A}_{k,i-1}, \mathcal{B}_{k,i-1})$, a has an E–successor b (the root of a copy of $H^{k+2,k}(\mathcal{A}_{k,i-1}, \mathcal{B}_{k,i-1})$) which has $k+2$ E–successors c such that $\mathcal{D}_c$ is isomorphic to $\mathcal{A}_{k,i-1}$, and k E–successors c' such that $\mathcal{D}_{c'}$ is isomorphic to $\mathcal{B}_{k,i-1}$. Hence, the induction hypothesis implies that b has exactly $k+2$ E–successors in the set $X(\mathcal{D})_j$. Since $p+k+2 = t+1$ and $d_{t+1} = +$, we conclude that $\mathcal{D} \models \chi[a, b, X(\mathcal{D})_j]$, and consequently $a \in X(\mathcal{D})_{j+1}$.
Assume next that $\mathcal{D}_a$ is isomorphic to $\mathcal{B}_{k,i}$ for some $i < j+1$. Then a has two successors b and b', and both $\mathcal{D}_b$ and $\mathcal{D}_{b'}$ are isomorphic to $H^{k+1,k+1}(\mathcal{A}_{k,i-1}, \mathcal{B}_{k,i-1})$. Thus, by the induction hypothesis, b (and b') has exactly $k+1$ E–successors in $X(\mathcal{D})_j$. Since $p+k+1 = t$ and $d_t = -$, we have $\mathcal{D} \not\models \chi[a, b, X(\mathcal{D})_j]$, and so $a \notin X(\mathcal{D})_{j+1}$.
Finally, if $\mathcal{D}_a$ is not isomorphic to either $\mathcal{A}_{k,i}$ or $\mathcal{B}_{k,i}$ for some $i < j+1$, then no E–successor c of any E–successor b of a is in $X(\mathcal{D})_j$, i.e., the second conjunct of $\chi[a, b, X(\mathcal{D})_j]$ fails for every b such that $(a, b) \in E^{\mathcal{D}}$. Hence, $a \notin X(\mathcal{D})_{j+1}$.

In particular, we have $u \in X(\mathcal{A}^{p,q}_{k,l})_\infty$, but $v \notin X(\mathcal{B}^{p,q}_{k,l})_\infty$, where u and v are the roots of the trees $\mathcal{A}_{k,l}$ and $\mathcal{B}_{k,l}$, respectively. Thus, $\mathcal{A}^{p,q}_{k,l} \models \varphi$ but $\mathcal{B}^{p,q}_{k,l} \not\models \varphi$, where φ is the sentence

$$\exists z(\exists y Ezy \wedge \forall y \neg Eyz \wedge [IFP_{x,X}\theta(x, X)]z).$$

It follows now from Proposition 21 that φ is not equivalent to any sentence of the logic $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$. Since $FO(Q_{IFP}, Q) \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$, we conclude that $IFP(Q) \not\leq FO(Q_{IFP}, Q)$. This completes the proof of Theorem 14.

5 Making quantifiers monotone

Let Q be a quantifier. We are going to define a monotone class Q^{mon} such that Q and Q^{mon} are first order definable from each other. This will enable us to translate our results about extensions of *IFP* into results on extensions of *LFP*.

Note that Definition 1 only states the Q–clause for the extension of a given logic. It does not say how to deal with restrictions in the other formation rules. In the case of *LFP* the notion of positivity has to be adapted to the case of extensions with quantifiers. We do this in the usual way, that is we define *LFP*(Q) only for the case of monotone Q, and we assume positivity, and negativity respectively, to be left unchanged by applications of Q. In this framework, one can rely on the known monotonicity properties of the least fixed point operator and hence extend many results on *LFP* to the presence of quantifiers. Especially, noting that the proof of the equivalence $LFP \equiv IFP$ does not depend on the type of occurring subformulas as long as the least fixed point operators are well defined, we see that the corresponding extension of the theorem of Gurevich and Shelah holds.

Remark 22: For each monotone quantifier Q, we have $LFP(Q) \equiv IFP(Q)$. $\square$

In order to extend *LFP* with an arbitrary quantifier Q, we pass to the following equivalent monotone version Q^{mon}. We get Q^{mon} by introducing new relation symbols for the complements of the relations in $\sigma = \{R_1 \ldots R_s\}$ (the vocabulary of Q, as in Definition 1).

Definition 23: For each $R_i \in \sigma$ let R_i^- be a relation symbol of the same arity as R_i. Let $\sigma^- := \sigma \cup \{R_i^- \mid 1 \leq i \leq s\}$. In the following $FO[\sigma^-]$–formulas, $\overline{x}_i$ denotes a tuple of distinct variables that matches the arity of R_i.

$$\varphi_\cap := \bigvee_{1 \leq i \leq s} \exists \overline{x}_i (R_i \overline{x}_i \wedge R_i^- \overline{x}_i),$$

$$\varphi_{PARTITION} := \bigwedge_{1 \leq i \leq s} \forall \overline{x}_i (R_i \overline{x}_i \leftrightarrow \neg R_i^- \overline{x}_i),$$

$$Q^{mon} := Mod(\varphi_\cap \vee (\varphi_{PARTITION} \wedge Q\overline{x}_1, \ldots, \overline{x}_s; \Phi_\sigma)), \text{ where}$$

$$\Phi_\sigma := R_1\overline{x}_1, \ldots, R_s\overline{x}_s.$$

Theorem 24:

(i) Q^{mon} is monotone.

(ii) Q and Q^{mon} are first order definable from each other, that is we have

$$Q \in FO(Q^{mon}) \text{ and } Q^{mon} \in FO(Q).$$

Consequently, $Q^r \in FO(Q^{mon})$ and $(Q^{mon})^r \in FO(Q)$ for each $r \geq 1$.

Proof: (i) Let $\mathcal{A} \in Q^{mon}$ and consider a σ^-–structure $\mathcal{A}'$ with $A' = A$ and $R_i \subset R_i'$ for all i. We have to show $\mathcal{A}' \in Q^{mon}$. If $\mathcal{A} \models \varphi_\cap$, then also $\mathcal{A}' \models \varphi_\cap$. If $\mathcal{A} \models \neg\varphi_\cap$ and $\mathcal{A}' \models \neg\varphi_\cap$, then we have $\mathcal{A} \models \varphi_{PARTITION}$ and the relations cannot have been properly extended. Hence $\mathcal{A}' = \mathcal{A}$ and the claim follows.
(ii) Clearly, Q^{mon} is definable in $FO(Q)$. On the other hand, we show that

$$Q = Mod(Q^{mon}\overline{x}_1, \ldots, \overline{x}_s, \overline{x}_1, \ldots, \overline{x}_s; (\Phi_\sigma, \neg\Phi_\sigma)),$$

where $\neg\Phi_\sigma$ is derived from Φ_σ by negating each formula in that tuple.

To see this, note first that a σ^-–structure $\mathcal{A}$ which is a model of $\varphi_{PARTITION}$ belongs to Q^{mon} if and only if its σ–reduct belongs to Q. Now, let $\mathcal{B}$ be a σ–structure, and set $\mathcal{A} := \mathcal{B}_{\Phi_\sigma, \neg\Phi_\sigma}$. Then $\mathcal{A} \models \varphi_{PARTITION}$ and $\mathcal{B}$ is the σ–reduct of $\mathcal{A}$. Hence $\mathcal{A} \in Q^{mon}$ iff $\mathcal{B} \in Q$, concluding the proof. $\square$

Corollary 25: For each unbounded simple cardinality quantifier Q, we have
$FO(Q_{LFP}, Q^{mon}) < LFP(Q^{mon})$.

Proof: The left side coincides with $FO(Q_{IFP}, Q)$, the right one with $IFP(Q^{mon})$ (by Remark 22) and hence with $IFP(Q)$. $\square$

6 The tree isomorphism problem

Recall that $\mathbf{C}$ denotes the class of all counting quantifiers $Q_i := \{(A; R) \mid |R| \geq i\}$. These Lindström versions of counting quantifiers do not increase the expressive power of TC or IFP since they are first order definable. But they do increase that of the infinitary logic $\mathcal{L}^\omega_{\infty\omega}$, even if their nesting is subject to a finite bound.

Another way to add counting to logics is given by the following *two sorted* approach. We explain it for the case of TC.

Associate with each structure $\mathcal{A}$ a copy of the set $\{0, \ldots, |A| - 1\}$ assumed to be disjoint from A. This second sort of elements is equipped with the natural order, $<$, of numbers. In addition to the individual variables $x, y, \ldots$ (running over the elements of A) there are now variables $i, j, \ldots$ running only over elements of the second sort. This gives rise to new atomic formulas $i < j$ and to quantifications like $\exists i\psi$. Similarly, in TC–operators one may use tuples mixed from both sorts. The counting comes in via the new clauses

$$\exists^{\geq i} x \psi(x)$$

meaning that there are at least i elements x such that $\psi(x)$. Note that this quantifier binds x whereas it introduces i as a free variable. We write $TC + COUNTING$ for the class of queries (over the first sort) definable in this way. In its fragment $TC + (l - COUNTING)$ the nesting of counting quantifiers $\exists^{\geq i} x$ is bounded by l. It should be clear how $\mathcal{L}^{\omega}_{\infty\omega} + (l - COUNTING)$ is defined.

Proposition 26: $TC + COUNTING \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{\omega}$.

Proof: We show that the inclusions $TC + (l - COUNTING) \leq \mathcal{L}^{\omega}_{\infty\omega} + (l - COUNTING) \leq \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C})^{l}$ hold for all $l \geq 1$. For the first inclusion one can argue in exactly the same way as without the presence of counting. For the second inclusion one can associate with each formula $\varphi(\bar{i}) \in \mathcal{L}^{\omega}_{\infty\omega} + (l - COUNTING)$ and each tuple of numbers $\bar{r}$ of the length of $\bar{i}$ another formula $\varphi_{\bar{r}} \in \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{l}$ such that the following holds: For all structures $\mathcal{A}$ of size at least $\max\{\bar{r}\}$ we have the equivalence

$$\mathcal{A} \models \varphi[\bar{r}] \text{ iff } \mathcal{A} \models \varphi_{\bar{r}}.$$

This is a well known argument and has been used, for instance, in [4], Chapter 7.4. The definition of $\varphi_{\bar{r}}$ is by induction on φ, simultaneously for all tuples $\bar{r}$. We sketch the most important steps. If $\varphi = i_1 < i_2$, we can set $\varphi_{\bar{r}} := TRUE$ for $r_1 < r_2$, or $\varphi_{\bar{r}} := FALSE$ for $r_1 \geq r_2$ respectively.

For $\varphi = \exists i_1 \psi$ note that there are formulas $\varphi_n \in \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{1}$ saying that the size of the domain is n. We can then restate φ as

$$\psi_{(r_2,\ldots)} := \bigvee_{n \geq 1} (\varphi_n \wedge \bigvee_{s_1 \leq n} \psi_{(s_1, r_2, \ldots)}).$$

The counting case, $\varphi = \exists^{\geq i_1} x \psi$, is now covered by

$$\psi_{\bar{r}} := \; Q_{r_1} x; \psi_{\bar{r}}.$$

On the right side, only the Lindström version of counting occurs and the nesting of counting quantifiers remains unchanged. □

Observe that the inclusion of the last proposition is strict. In fact, one can easily find nonrecursive cardinality statements in $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{C}^*)^{1}$.

Although the two sorted counting is a considerable increase in expressive power for *TC* there are still *LOGSPACE* computable problems which are not definable in $TC + COUNTING$. This has recently been proved by Etessami and Immerman who studied the *tree isomorphism problem* [5]. It can be presented, with binary E and unary P, as the class

$$Q_{TI} := \{(A; E, P) \mid (P; E \cap P^2) \text{ and } (A \backslash P; E \cap (A \setminus P)^2) \text{ are isomorphic trees}\}.$$

Theorem 27 [5]: $Q_{TI} \notin TC + COUNTING$. □

Using the trees $\mathcal{A}_{k,l}$ and $\mathcal{B}_{k,l}$ constructed in Section 4, we are able prove that tree isomorphism cannot even be defined in infinitary logic with monadic quantifiers of bounded nesting and with a finite bound on the number of variables.

Theorem 28: $Q_{TI} \notin \mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$.

Proof: Given two τ–structures $\mathcal{A}$ and $\mathcal{B}$, their *disjoint sum* $\mathcal{A} \oplus \mathcal{B}$ is a $\tau \cup \{P\}$–structure such that the universe of $\mathcal{A} \oplus \mathcal{B}$ is the disjoint union of A and B, $P^{\mathcal{A}\oplus\mathcal{B}} = A$, and $R^{\mathcal{A}\oplus\mathcal{B}} = R^{\mathcal{A}} \cup R^{\mathcal{B}}$ for each relation symbol $R \in \tau$. Using the k,l–game, it is easy to show that $\equiv_{k,l}$–equivalence is preserved by disjoint sums: if $\mathcal{A} \equiv_{k,l} \mathcal{A}'$ and $\mathcal{B} \equiv_{k,l} \mathcal{B}'$, then $\mathcal{A} \oplus \mathcal{B} \equiv_{k,l} \mathcal{A}' \oplus \mathcal{B}'$.

In particular, by Proposition 19 we have $\mathcal{A}_{k,l+1} \oplus \mathcal{A}_{k,l+1} \equiv_{k,l} \mathcal{A}_{k,l+1} \oplus \mathcal{B}_{k,l+1}$. Since $\mathcal{A}_{k,l+1} \oplus \mathcal{A}_{k,l+1} \in Q_{TI}$ and $\mathcal{A}_{k,l+1} \oplus \mathcal{B}_{k,l+1} \notin Q_{TI}$, we conclude that Q_{TI} is not definable in $\mathcal{L}^{\omega}_{\infty\omega}(\mathbf{Q_1})^{\omega}$. □

By virtue of Proposition 26, Theorem 28 implies the result of Etessami and Immerman. As a matter of fact, it implies that Q_{TI} is not definable in $FO(Q)+COUNTING$ for any $\mathcal{L}^{\omega}_{\infty\omega}$–definable quantifier Q, including all quantifiers Q_{IFP} and Q_{PFP} such that $FO(Q_{IFP}) \equiv IFP$ and $FO(Q_{PFP}) \equiv PFP$.

References

[1] S. Abiteboul and V. Vianu. Fixpoint extensions of first order logic and datalog-like languages. In *Proc. 4th IEEE Symp. on Logic in Computer Science*, pages 71–79, 1989.

[2] E. Dahlhaus. Skolem normal forms concerning the least fixpoint. In E. Börger, editor, *Computation theory and Logic, Lecture Notes in Computer Science 270*, pages 101–106. Springer-Verlag, 1987.

[3] H.-D. Ebbinghaus. Extended logics: The general framework. In J. Barwise and S. Feferman, editors, *Model–Theoretic Logics*, pages 25–76. Springer-Verlag, 1985.

[4] H.-D. Ebbinghaus and J. Flum. *Finite Model Theory.* Springer-Verlag, 1995.

[5] K. Etessami and N. Immerman. Tree canonization and transitive closure. In *Proc. 10th IEEE Symp. on Logic in Computer Science*, pages 331–341, 1995.

[6] M. Grohe. Complete problems for fixed–point logics. *Journal of Symbolic Logic*, 60:517–527, 1995.

[7] Y. Gurevich and S. Shelah. Fixed point extensions of first–order logic. *Annals of Pure and Applied Logic*, 32:265–280, 1986.

[8] L. Hella. Definability hierarchies of generalized quantifiers. *Annals of Pure and Applied Logic*, 43:235–271, 1989.

[9] L. Hella. Logical hierarchies in PTIME. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, 1992. Full version to appear in *Information and Computation.*

[10] L. Hella. Fixpoint logic vs generalized quantifiers, 1994. Manuscript.

[11] L. Hella, J. Väänänen and D. Westerståhl. Definability of polyadic lifts of generalized quantifiers, 1995. Submitted to *Journal of Logic, Language and Information.*

[12] H. Imhof. *Fixed Point Logics and Generalized Quantifiers in Descriptive Complexity.* PhD thesis, Albert-Ludwigs Universität Freiburg, 1996.

[13] H. Imhof. Fixed point logics, generalized quantifiers and oracles, 1995. To appear in *Journal of Logic and Computation.*

[14] N. Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[15] Ph. G. Kolaitis and J. Väänänen. Generalized quantifiers and pebble games on finite structures. *Annals of Pure and Applied Logic*, 74:23–75, 1995.

[16] Ph. G. Kolaitis and M. Y. Vardi. Infinitary logics and 0–1 laws. *Information and Computation*, 98:258–294, 1992.

[17] Ph. G. Kolaitis and M. Y. Vardi. Fixpoint logic vs. infinitary logic in finite-model theory. In *Proc. 7th IEEE Symp. on Logic in Computer Science*, 1992.

[18] P. Lindström. First order predicate logic with generalized quantifiers. *Theoria*, 32:186–195, 1966.

[19] M. Y. Vardi. The complexity of relational query languages. In *Proc. of the 14th ACM symposium on the theory of computing*, pages 137–146, 1982.