Orderings and PTIME – On a Conjecture by Makowsky

Matthias Ruhl
MIT Laboratory for Computer Science
Cambridge MA, 02139, USA
ruhl@theory.lcs.mit.edu

August 24, 2001

Abstract

Immerman and Vardi showed in the early 1980s that the logic LFP, augmented with an ordering, can express all PTIME-decidable properties. In 1997, Makowsky conjectured that LFP+$\mathcal{A}$, i.e. LFP augmented with an arbitrary relation $\mathcal{A}$, captures PTIME if and only if a total ordering is parametrically definable on $\mathcal{A}$ using LFP. This would establish that LFP's computational power crucially depends on the input being ordered.

In this paper, we disprove Makowsky's conjecture by giving a class of relations $\mathcal{A}$ such that LFP+$\mathcal{A}$ captures PTIME, yet no ordering with more than $O(\sqrt{n})$ elements can be defined on $\mathcal{A}$.

1 Introduction

About 20 years ago, Immerman [Imm86] and Vardi [Var82] showed that the logic **LFP** captures the complexity class **PTIME** on ordered structures. That is, sentences from **LFP** can be decided in polynomial time, and every **PTIME**-decidable property on ordered structures can be expressed as a sentence in **LFP**.

In turns out that this characterization does not hold if no ordering is present. In fact, even Parity cannot be expressed in **LFP** without an ordering [CH82].

It has been a long open question what the exact role of orderings in capturing **PTIME** is, and whether adding weaker structures to the input could still allow the capturing of **PTIME** with **LFP**. By adding structure to the input we mean that given a class of *tau*-structures $\mathcal{A}$ (containing one structure for each finite cardinality), for example total orders, we consider the expressiveness of **LFP** on $(\sigma \cup \tau)$-structures $\mathfrak{B}$ whose τ-relations are isomorphic to an element of $\mathcal{A}$. We call these $\mathfrak{B}$ structures with built-in relations from $\mathcal{A}$, or simply speak about the expressiveness of **LFP**+$\mathcal{A}$.

In 1997, Eric Rosen proved the following result about a particular class of built-in relations [Mak97, Ros98], which is based on earlier work by Hella [Hel96]. An order of equivalence classes is a total order on the equivalence classes of an equivalence relation.

Theorem 1 (Rosen 1997)
If $\mathcal{A}$ is a class of orders of equivalence classes, then **LFP**+$\mathcal{A}$ *captures* **PTIME** *iff there is some c such that the number of equivalence classes in all structures $\mathfrak{A} \in \mathcal{A}$ are at least $|\mathfrak{A}| - c$.*

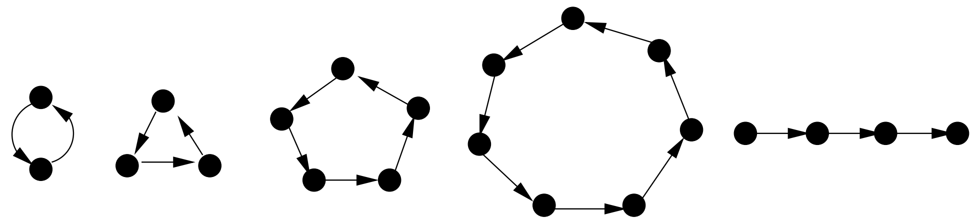

Figure 1: An example: A_{21}

This result shows that for these structures $\mathcal{A}$, capturing **PTIME** is equivalent to being able to parametrically define an ordering on $\mathcal{A}$. By "parametrically defining an ordering", we mean that there is a formula $\varphi(x, y, z_1, \ldots, z_k)$ ($k \geq 0$) in **LFP**, such that for every structure $\mathfrak{A} \in \mathcal{A}$, there exist elements $c_1, \ldots, c_k \in A$ (where A is the universe of $\mathfrak{A}$), such that $\{(x, y) \mid \mathfrak{A} \models \varphi(x, y, c_1, \ldots, c_k)\}$ is a total ordering on A.

Based on Rosen's result, Johann Makowsky conjectured in 1997 that this relationship between capturing **PTIME** and being able to define an ordering is true for all built-in relations [Mak97].

Conjecture 2 (Makowsky 1997)
Let $\mathcal{A}$ be a class of finite structures, containing, up to isomorphism, one structure of each cardinality. Then the following two statements are equivalent:

(i) **LFP**+$\mathcal{A}$ *captures* **PTIME**.

(ii) One can parametrically define an ordering on the structures in $\mathcal{A}$ with a formula in **LFP**.

Note that due to Immerman and Vardi's result, (ii) implies (i).

To be precise, Makowsky made a more general conjecture than stated above; he allowed for classes $\mathcal{A}$ with more than one (non-isomorphic) structure per cardinality. To avoid defining what capturing means in that case, we just consider this more restrictive version of the conjecture.

In this paper we show that Makowsky's conjecture is false by showing that (i) does not imply (ii). Showing that the more restrictive version of Makowsky's conjecture is false obviously implies that the general version is false as well. For our proof, we give a class $\mathcal{A}$ of structures such that one cannot define an ordering on them in **LFP**, but **LFP**+$\mathcal{A}$ still captures **PTIME**.

2 The Counterexample

In the following, let p_i be the i-th prime number ($p_1 = 2$, $p_2 = 3$, $p_3 = 5$, ...). We now define the structure class $\mathcal{A}$ that will be the counterexample to Makowsky's conjecture.

Definition 3 ($\mathcal{A}$)
Let $\mathcal{A}$ be the set $\{A_1, A_2, A_3, \ldots\}$, where $A_n = (\{1, 2, \ldots, n\}, E_n)$. The binary relation E_n is defined as follows. Let k be the integer with the property that $p_1 + p_2 + \cdots + p_k \leq n$, but $p_1 + p_2 + \cdots + p_{k+1} > n$. Then E_n is the disjoint union of k directed cycles of sizes $p_1, p_2, \ldots, p_k$, and a directed line on the remaining $n - \sum_{i=1}^{k} p_i$ vertices. (See Figure 1 for an example.)

Lemma 4
There is no **LFP**-*formula that (even parametrically) defines a total order on all structures in* $\mathcal{A}$.

Proof: Since **LFP** cannot distinguish the elements of a cycle, one has to use a parameter to define an order on a cycle, and therefore needs at least as many parameters in the formula as there are cycles in the graph. But since the number of cycles in the structures in $\mathcal{A}$ is unbounded, there exists no **LFP** formula that parametrically defines an ordering on all the structures in $\mathcal{A}$. $\square$

3 Capturing PTIME

3.1 Interpretations

While it is not possible to define a total order on the structures in $\mathcal{A}$, one can use built-in relations from $\mathcal{A}$ to define (in **LFP**) an ordered copy of a given structure. This is enough to show that we capture **PTIME** by **LFP**+$\mathcal{A}$, since we can now just "compute" on the ordered copy.

The definition of the ordered copy is done by means of a logical interpretation (cf. [EF99]). Assume that we are considering structures $\mathfrak{B}$ of some signature $\sigma \cup \{E\}$, where E is the built-in relation from $\mathcal{A}$. We will then give **LFP**$[\sigma \cup \{E\}]$ formulas $\varphi_{\text{univ}}(x_1, x_2, c)$, $\varphi_{=}(x_1, x_2, y_1, y_2, c)$, $\varphi_{<}(x_1, x_2, y_1, y_2, c)$, and $\varphi_R(x_1, \ldots, x_{2r_R}, c)$ for all relation symbols $R \in \sigma$ with arity r_R. The constant c will be fixed appropriately later. This set of formulas implicitly defines a structure with signature $\sigma \cup \{<\}$ as follows:

- The universe of the new structure is[1]

  $$U = \{(x_1, x_2) \mid \mathfrak{B} \models \varphi_{\text{univ}}(x_1, x_2, c)\} \ / \ \{((x_1, x_2), (y_1, y_2)) \mid \mathfrak{B} \models \varphi_{=}(x_1, x_2, y_1, y_2, c)\},$$

  where $\varphi_{=}$ defines an equivalence relation on $B \times B$, where B is the universe of $\mathfrak{B}$. The set U is therefore a set of equivalence classes on a subset of $B \times B$.

- The binary relation $\{([(x_1, x_2)], [(y_1, y_2)]) \in U \mid \mathfrak{B} \models \varphi_{<}(x_1, x_2, y_1, y_2, c)\}$ is a total ordering on U.

- The φ_R similarly define r-ary relations on U.

The following is the main result of this paper.

Lemma 5
Let σ be a finite signature. Then there exists an **LFP**+$\mathcal{A}$-*interpretation π (in the above sense) such that for any σ-structure $\mathfrak{B}$, the structure defined under the interpretation $\mathfrak{B}^\pi$ is isomorphic to $\mathfrak{B}$ and exhibits a total ordering.*

Actually, by being careful in the construction, one can achieve a **DTC**+$\mathcal{A}$ interpretation that achieves the same thing.

Before we prove the lemma, let us briefly consider why this is sufficient to prove that **LFP**+$\mathcal{A}$ captures **PTIME**. We only have to show that for every sentence φ in **LFP**+$<$ there is a sentence φ' in **LFP**+$\mathcal{A}$ such that φ is valid on a structure with built-in order iff φ' is valid on the same structure with a built-in relation from $\mathcal{A}$ instead. Given our interpretation from Lemma 5, we can define φ' from φ inductively, replacing every variable x by a pair of variables (x_1, x_2):

[1] By writing X/Y for a set X and an equivalence relation Y on a superset of X, we denote the set of equivalence classes in $(X \times X) \cap Y$.

- $[x = y]' = \varphi_=(x_1, x_2, y_1, y_2)$
- $[x < y]' = \varphi_<(x_1, x_2, y_1, y_2)$
- $[\varphi \wedge \psi]' = \varphi' \wedge \psi'$
- $[\exists x \varphi]' = \exists x_1 \exists x_2 \ \varphi_{\text{univ}}(x_1, x_2) \wedge \varphi'$

More details can be found in [EF99], in particular the more involved inductive step for the LFP-operator.

3.2 Constructing an Ordered Copy

Proof (Lemma 5): For simplicity's sake, we assume in the following that all elements of B are part of a cycle in E. If that were not the case, the additional elements would be totally ordered, which makes the job of defining an ordered copy even easier.

First, we observe that we can define an ordering on the cycles in $\mathfrak{A}$, since one can express in **LFP** that one cycle is bigger than another cycle.

For the following let c be any element of the largest cycle in the relation E. This allows us to order the largest cycle (by, for example, letting c be the minimal element, and following the natural order of the cycle). We will identify this order with the numbers $\{1, 2, \ldots, p_k\}$.

The universe of our ordered copy will be the following set:

$$\begin{aligned}
&(C_2, 1), (C_2, 2),\\
&(C_3, 1), (C_3, 2), (C_3, 3),\\
&(C_5, 1), (C_5, 2), (C_5, 3), (C_5, 4), (C_5, 5),\\
&\ldots,\\
&(C_{p_k}, 1), \ldots, (C_{p_k}, p_k)
\end{aligned}$$

Here the first element of every pair is an equivalence class containing all elements of the given cycle (C_p denotes the cycle of size p). This is accomplished by setting

- $\varphi_{\text{univ}}(x_1, x_2, c)$ = "x_2 is an element of c's cycle" $\wedge$ "the number of elements of x_1's cycle is more than the number of edges on the path from c to x_2"
- $\varphi_=(x_1, x_2, y_1, y_2, c)$ = "x_1 and y_1 are on the same cycle"
- $\varphi_<(x_1, x_2, y_1, y_2, c)$ = "y_1's cycle is larger than x_1's cycle" $\vee$ ("x_1 and y_1 are on the same cycle" $\wedge$ "y_2 is further from c than x_2")

To define the φ_R's we will try to create a mapping between each cycle C_p and the elements $(C_p, 1), (C_p, 2), \ldots, (C_p, p)$ of the interpreted structure. It will turn out that either we can use the structure present in $\mathfrak{B}$ to create a one-to-one mapping between the two, or that all elements on the cycle "behave the same" under all relations, in which case we can define an ordered copy without creating any particular mapping. This follows from the following lemma.

Lemma 6
The automorphism group $Aut(\mathfrak{B})$ of $\mathfrak{B}$ is of the form

$$Aut(\mathfrak{B}) = G_2 \times G_3 \times G_5 \times \cdots \times G_{p_k},$$

where each $G_p \in \{0, \mathbb{Z}_p\}$ (by 0 we denote the trivial group).

Proof: First note that the automorphism group of the cycles (relation E) by themselves is $G = \mathbb{Z}_2 \times \mathbb{Z}_3 \times \cdots \times \mathbb{Z}_{p_k}$. Since $Aut(\mathfrak{B})$ is a subgroup of G, it is certainly enough to show that all subgroups of G are of the form stated in the lemma. As G is cyclic (generated by the element $(1,1,\dots,1)$), any subgroup of G is also cyclic. So consider some subgroup generated by the element $g = (g_1, g_2, \dots, g_k)$.

For each i such that $g_i \neq 0$, by the Chinese Remainder Theorem there is some integer x such that $x \cdot g_i \equiv 1\ (p_i)$ and $x \cdot g_j \equiv 0\ (p_j)$ for all $j \neq i$. This shows that the element $g^x = (0, 0, \dots, 0, 1, 0, \dots, 0)$ (with the 1 at the i-th position) is in the subgroup. Since that is true for all i with $g_i \neq 0$ it is not hard to see that the subgroup generated by these elements is equal to $G_{p_1} \times G_{p_2} \times \cdots \times G_{p_k}$, where $G_{p_i} = 0$ if $g_i = 0$ and $G_{p_i} = \mathbb{Z}_{p_i}$ if $g_i \neq 0$. This is because clearly we do not get more than the subgroup, but also at least as much, since in particular we obtain the generator back. This shows that the subgroup is of the desired form. □

3.3 Ordering Cycles

The formulas φ_R will be of the following form:

$$\begin{aligned}\varphi_R(x_1, x_2, \dots, x_{2r_R}, c) = \exists y_1 y_2 \dots y_{r_R} \text{“}y_1 = f(x_1, x_2)\text{”} \wedge \text{“}y_2 = f(x_3, x_4)\text{”} \wedge \dots \\ \cdots \wedge \text{“}y_{r_R} = f(x_{2r_R - 1}, x_{2r_R})\text{”} \wedge R y_1 y_2 \dots y_{r_R}\end{aligned}$$

Here f is an isomorphism (defined in **LFP**+$\mathcal{A}$ between the universe U of the ordered copy and the universe B of the structure $\mathfrak{B}$. The function f actually depends on the values of the parameters $(x_i)_{1 \le i \le r_R}$, but only in a way that does not impact the validity of whether $R y_1 y_2 \dots y_{r_R}$ holds or not.

To define f, we proceed as follows:

1. First, we define a unary relation F, which contains exactly one element of each cycle C_p, for which $G_p = 0$, and no other elements.

2. Given the values of the parameters $x_1, x_2, \dots, x_{r_R}$ we augment F such that it contain a single element for all cycles on which the element $\{x_{2i-1} \mid 1 \le i \le r_R\}$ are (recall the structure of the universe of our ordered copy: the first element of each pair denotes a cycle, while the second encodes an index on that cycle). This is done by setting

   $F' := F \cup \{x_{2i-1} \mid$ “x_{2i-1} is on a different cycle than $x_{2j-1} \forall j < i$” $\wedge F \cap$ “x_{2i-1}'s cycle” $= \emptyset\}$

3. Now we define $f(a_1, a_2)$. Let z be the unique element in F' on a_1's cycle. Then $f(a_1, a_2)$ is equal to the element which is as many steps from z (on z's cycle) as a_2 is from c.

The function f maps each of element of the form (C_p, i) to an element of C_p. For cycles with $G_p = 0$ this mapping is fixed based on the elements of F. For all other cycles, since $G_p = \mathbb{Z}_p$ any assignment that preserves the relative positions of elements on the same cycle will lead to an isomorphism. This is why while the extension of F to F' depend on the values of the x_i, the validity of $Ry_1y_2 \dots y_{r_R}$ does not depend on them.

So to complete the construction, all we have to do is to define the relation F, which singles out one element per cycle C_p with $G_p = 0$. For the following, let r be the maximum arity of any relation in σ. This means that r is a constant dependent only on σ.

Consider some cycle C_p for which the factor G_p in $Aut(\mathfrak{B})$ is the trivial group. In particular, simply rotating the cycle C_p is not an automorphism. So there exist some relations involving the elements of this cycle that do not hold after rotating the cycle. In particular there must be some relation $R \in \sigma$ and a tuple $(x_1, \dots, x_{r_R})$ of elements such that $Rx_1 \dots x_{r_R}$ before the rotation, but $\neg R\rho(x_1) \dots \rho(x_{r_R})$ after the rotation ρ. So the rotation is already not an automorphism on the set of cycles containing the r_R elements x_i. Let us make that first observation.

Fact 7
For every cycle C_p with $G_p = 0$, there is a set of r cycles (which includes C_p), such that even restricted to this set of cycles, any non-trivial rotation of C_p is not an automorphism.

This fact that we can detect a non-automorphism "locally" will enable us to distinguish the elements of C_p within **LFP**. So let's consider such a set of r cycles, containing C_p.

Suppose we fix one element on each of these r cycles. By considering these elements to be the minimal elements on their cycles, we can construct a total order on these r cycles. This allows us to write down the values of the relations on these cycles as a 0-1-sequence, by stepping through the arguments in lexicographic order, and recording a 1 if the relation holds, and 0 if it does not hold. Obviously, this sequence depends on the particular ordering, and therefore the elements we fixed on each cycle. In particular, since picking a different element on C_p corresponds to a rotation of C_p, a different choice of the element on C_p (no matter whether the other elements changes as well or not) leads to a different sequence.

In **LFP**, we can determine which choice of r elements on our cycles leads to the lexicographically smallest 0-1-sequence[2]. The choice of elements on the cycles which leads to this minimum need not be unique. But all choices that lead to the minimum sequence have the same fixed element on C_p. This is because different choices of elements on C_p lead to different sequences. This enables us to uniquely identify this element on C_p, and add it to our relation F.

Initially, we do not know which G_p are the trivial group, and which set of r cycles exhibits that fact for each p. To make sure that we find all such p, we simply step though all possible subsets of r cycles (there are only polynomially many, since r is constant), and apply the above minimum sequence finding algorithm. We maintain the unary relation F, initially empty, which contains the fixed elements. When we identify a single element on some cycle C_p, and no element of that cycle is yet in F, we add that element to F. After having considered all sets of r cycles, for each p with $G_p = 0$ there will be a single element in $F \cap C_p$, while $F \cap C_p = \emptyset$ for all other p.

This enables us to construct φ_R as above, and concludes the proof. □

[2] This is because the cycles are completely ordered, so we can express anything computable in **PTIME**, in particular which of the (at most) n^r choices of elements on the r cycles leads to the minimal 0-1-sequence.

References

[CH82] Ashok Chandra and David Harel. Structure and complexity of relational queries. *Journal of Computer and System Sciences (JCSS)*, 25(1):99–128, 1982.

[EF99] Heinz-Dieter Ebbinghaus and Jörg Flum. *Finite Model Theory.* Springer Verlag, 1999.

[Hel96] Lauri Hella. Logical hierarchies in ptime. *Information and Computation*, 129:1–19, 1996.

[Imm86] Neil Immerman. Relational queries computable in polynomial time. *Information and Control*, 68:86–104, 1986.

[Mak97] Johann Makowsky. Invariant definability. In *Computational Logic and Proof Theory, Proceedings of the 5th Kurt Gödel Colloquium, KGC'97, LNCS vol. 1289*, pages 186–202, 1997.

[Ros98] Eric Rosen. Personal communication, 1998.

[Var82] Moshe Y. Vardi. The complexity of relational query languages. In *Proc. 14th ACM Symp. on Theory of Computing*, 1982.