Epsilon-Logic is More Expressive than First-Order Logic over Finite Structures

Martin Otto

revised May 99

Abstract

There are properties of finite structures that are expressible with the use of Hilbert's ϵ-operator in a manner that does not depend on the actual interpretation for ϵ-terms, but not expressible in plain first-order. This observation strengthens a corresponding result of Gurevich, concerning the invariant use of an auxiliary ordering in first-order logic over finite structures. The present result also implies that certain non-deterministic choice constructs, which have been considered in database theory, properly enhance the expressive power of first-order logic even as far as deterministic queries are concerned, thereby answering a question raised by Abiteboul and Vianu.

Introduction

Hilbert's ϵ-operator extends first-order by means of a choice construct, which allows to select elements from non-empty definable subsets of the universe.

Epsilon-terms and epsilon-formulae. With each formula φ associate a term $\epsilon_x\varphi$ and put $\mathrm{free}(\epsilon_x\varphi) = \mathrm{free}(\varphi) \setminus \{x\}$. Let $\mathrm{FO}[\epsilon]$ be the extension of FO that is closed under the formation and use of ϵ-terms along with the usual first-order formula formation rules. In general, the formulae of $\mathrm{FO}[\epsilon]$ require so-called ϵ-structures for their semantics, namely structures which provide an interpretation for ϵ-terms.

Epsilon-structures. An ϵ-structure of signature τ consists of an ordinary τ-structure $\mathfrak{A} = (A, \dots)$ together with a *choice function* on the power set $\mathcal{P}(A)$ of the universe A of $\mathfrak{A}$:

$$\begin{aligned} \epsilon\colon \mathcal{P}(A) &\longrightarrow A \\ s &\longmapsto \epsilon(s) \quad \text{such that } \epsilon(s) \in s \text{ whenever } s \neq \emptyset. \end{aligned}$$

$\mathrm{FO}[\epsilon]$-terms are interpreted over ϵ-structures in the obvious way. If $\varphi = \varphi(\overline{y}, x)$, and for an assignment of elements $\overline{c}$ to the variables $\overline{y}$:

$$(\epsilon_x\varphi)^{\mathfrak{A},\epsilon}[\overline{c}] = \epsilon\big(\{a \in A : (\mathfrak{A}, \epsilon) \models \varphi[\overline{c}, a]\}\big).$$

Ultimately one is mostly interested, however, in those ϵ-formulae whose semantics does not depend on the actual interpretation of the choice function.

Definition 1 *A formula $\varphi(\overline{x}) \in \mathrm{FO}[\epsilon]$ is called ϵ-invariant (over some class of structures $\mathcal{K}$), if for any structure $\mathfrak{A}$ (from $\mathcal{K}$), and any two interpretations ϵ_1 and ϵ_2 as choice functions on $\mathcal{P}(A)$, and for all $\overline{a}$ from A: $(\mathfrak{A}, \epsilon_1) \models \varphi[\overline{a}]$ if and only if $(\mathfrak{A}, \epsilon_2) \models \varphi[\overline{a}]$.*

An easy exercise based on Craig's interpolation theorem for first-order logic (compare e.g. [9]) shows that, as far as the class of *all* structures is concerned, ϵ-invariance (like any other form of invariance under FO-axiomatizable expansions and extensions) is too strong a requirement to give rise to anything new:

Fact 2 *Any formula of $\mathrm{FO}[\epsilon]$ that is ϵ-invariant over the class of all structures is logically equivalent with a plain first-order formula.*

Various applications of choice constructs, in particular in database theory, as well as a general theoretical interest in finite model theory suggest to consider $\mathrm{FO}[\epsilon]$ over finite structures and under the corresponding, weaker requirement of invariance.

Definition 3 *Let $\mathrm{FO}[\epsilon]_{inv} \subseteq \mathrm{FO}[\epsilon]$ be the fragment of $\mathrm{FO}[\epsilon]$ consisting of those formulae that are ϵ-invariant over the class of all finite structures.*

As the interpolation theorem for first-order logic is known to fail in restriction to finite structures (compare [10]), there is no a priori reason that $\mathrm{FO}[\epsilon]_{\mathrm{inv}}$ should be a trivial extension of FO, and indeed the purpose of this note is to show that it is not.

Choice and order. Another extension of first-order logic over finite structures, which is similar in spirit, involves the expansion of structures by means of an arbitrary linear ordering of the universe, rather than by a choice function. Let $\mathrm{FO}[<]$ be ordinary first-order logic, where a new binary $<$ is added to the underlying vocabulary. A formula of $\mathrm{FO}[<]$ is $<$-invariant over $\mathcal{K}$ if its semantics is independent of the interpretation of $<$ as long as $<$ is interpreted as a linear ordering of the universe. Again first-order interpolation implies that $\mathrm{FO}[<]$-formulae that are $<$-invariant over *all* structures are equivalent with FO-formulae not involving $<$ at all. Let $\mathrm{FO}[<]_{\mathrm{inv}}$ stand for the fragment of $\mathrm{FO}[<]$ consisting of those formulae that are $<$-invariant over the class of all finite structures. Gurevich has shown that $\mathrm{FO} \subsetneq \mathrm{FO}[<]_{\mathrm{inv}}$, i.e. there is a property of finite structures (namely e.g. the evenness of the number of atoms in a finite Boolean algebra) that is first-order expressible in an $<$-invariant manner with the use of an auxiliary order, but not without (see [10], and Exercise number 17.27 in [1]).

Now $\mathrm{FO}[\epsilon]_{\mathrm{inv}}$ actually is a sublogic[1] of $\mathrm{FO}[<]_{\mathrm{inv}}$. This follows easily from the observation that linearly ordered finite domains provide a first-order de-

[1] According to most definitions, neither ϵ-invariant nor $<$-invariant first-order would qualify as *logics*, as the invariance properties involved in their definitions are undecidable. It is actually an interesting open question, whether there are recursive subsystems of $\mathrm{FO}[\epsilon]$ or of $\mathrm{FO}[<]$ that represent exactly all and only the invariantly definable properties over finite structures.

finable choice function, namely the one that always picks $<$-least elements. If the linear order $<$ is employed to simulate an invariant use of the ϵ-operator, then clearly the resulting FO[$<$]-formula is $<$-invariant. We thus strengthen Gurevich's result in showing that even FO $\subsetneq$ FO$[\epsilon]_{\mathrm{inv}}$. This still leaves open the status of the other inclusion: it would be nice to know whether indeed FO $\subsetneq$ FO$[\epsilon]_{\mathrm{inv}}$ $\subsetneq$ FO$[<]_{\mathrm{inv}}$, as one might expect.

Theorem 4 FO $\subsetneq$ FO$[\epsilon]_{inv}$ *over finite structures, i.e. there is a property of finite structures which is not* FO*-definable but definable in an* ϵ*-invariant manner with the use of the* ϵ*-operator.*

An application to a question in database theory

Non-deterministic choice constructs have been added to first-order logic in the context of relational database query languages, e.g. in the form of a *witness*-construct [1]. Apart from the usual first-order operations, the resulting language FO[*witness*] provides an operation *witness*, which applies to some $(n+1)$-ary predicate defined by a formula $\varphi(\overline{x}, y)$ and non-deterministically returns *some* $(n+1)$-ary predicate $R = witness_y\varphi$ satisfying

$$\forall\overline{x}\big(\exists y\varphi(\overline{x},y) \leftrightarrow \exists y R\overline{x}y\big) \,\wedge\, \forall\overline{x}\exists^{\leqslant 1} y R\overline{x}y.$$

As the name suggests, $witness_y\varphi$ selects unique y-witnesses or, equivalently, provides a choice function for the family $\{y \mid \varphi(\overline{x},y)\}_{\overline{x}}$ – potentially depending on the the formula φ and the parameters $\overline{x}$ in φ rather than just the set defined by $\varphi(\overline{x},\cdot)$. Semantically, the choice provided by the *witness* operation also differs from the choice provided by FO[ϵ] in being truly non-deterministic and independent in the sense that

(i) there is no pre-set interpretation for the *witness*-choices.

(ii) a priori even two consecutive *witness*-applications to the same φ need not return the same predicates.

The second, and somewhat disturbing feature of plain FO[*witness*] is remedied in a variant FO[*witness*]$^+$. This is the variant of [1], denoted FO$^+$+W in [2, 3], and also the variant to be considered here. For FO[*witness*]$^+$, FO itself is regarded as a query language that assigns relational values to new predicates. Correspondingly, the *witness* operation can be used in *assignments* of the form

$$R := witness_y\varphi(\overline{x},y),$$

which effectively serve to freeze one (non-deterministically chosen) witnessing predicate for φ, which can elsewhere be referred to by its name R. [2]

[2]While this extension of FO is very natural if FO is regarded as a query language in a relational algebra style, it should be noted that the possibility to locally freeze *witness*-choices during the evaluation of a formula or program essentially bypasses some difficulties inherent to the idea of truly independent choices. The reader should compare the analysis by Blass and Gurevich [5], where a form of truly independent choice is studied in contrast with ϵ-choice.

Of course one is really interested in uses of the witnessing mechanism that are independent of the actual choices made. It is customary in this context to call an FO[*witness*]- or FO[*witness*]$^+$-query *deterministic* (or *functional*) if its semantics is invariant under all admissible variations on the predicates returned for *witness*-applications (and over the class of all finite structures, i.e. for all database instances). The following theorem will be an immediate corollary to the proof of Theorem 4 to be given below.

Theorem 5 *Deterministic* FO[*witness*]$^+$ *is strictly more expressive than* FO *over finite relational structures (databases).*

Theorem 5 explicitly answers a question – or actually refutes one part of a conjecture – of Abiteboul and Vianu [3], and Abiteboul, Simon, and Vianu [2]. The other part of their conjecture, concerning the weaker variant FO[witness] without the freezing option, remains open.

It is clear that deterministic FO[*witness*]$^+$ is contained in FO$[\epsilon]_{\mathrm{inv}}$, since the *witness*-operation can be simulated in FO$[\epsilon]$ according to

$$witness_y \varphi(\overline{x}, y) := \big\{ (\overline{x}, \epsilon_y \varphi(\overline{x}, y)) \;\big|\; \exists y \varphi(\overline{x}, y) \big\}.$$

Note that there seems to be no straightforward simulation of ϵ-choice in FO[*witness*]$^+$, which would prove the converse inclusion. This is because of the fact that ϵ-choice from a definable set may depend only on that set, and not on its definition. In other words, the following is a validity of FO$[\epsilon]$, for all pairs of formulae $\varphi(\overline{x}, y), \psi(\overline{z}, y)$:

$$\forall \overline{x} \forall \overline{z} \big(\forall y (\varphi(\overline{x}, y) \leftrightarrow \psi(\overline{z}, y)) \longrightarrow \epsilon_y \varphi = \epsilon_y \psi \big).$$

To achieve a similar robustness by means of a *witness* assignment, one would like to single out one definition for each definable set over the given structure and then use a witness predicate for that particular definition. This, however, cannot be done by first-order means.

It thus remains open whether the containment of deterministic FO[*witness*]$^+$ in FO$[\epsilon]_{\mathrm{inv}}$ is strict. Similarly, the status of the obvious two separate inclusions, of FO in deterministic FO[*witness*] and of deterministic FO[*witness*] in deterministic FO[*witness*]$^+$, remains open: these may well both be strict; by Theorem 5 at least one of them has to be strict.

We also remark that in the presence of suitable recursion mechanisms the use of an order, of the ϵ-operator, and of the witness constructs all give rise to the same level of expressiveness. In particular, when added to least fixed-point logic all of these precisely allow to express all PTIME queries over finite structures in an invariant way (see e.g. the textbooks [10, 12] about capturing PTIME in the presence of order). The reason for the claimed equivalence is, of course, that successive choice (of any kind) can be used to define an ordering.

The separating example

The rest of this note is devoted to the proof of Theorems 4 and 5 through the analysis of a concrete example. The overall structure of the example can be used as a mold for certain variations, which also lead to a strengthening of the main theorem that will be briefly discussed at the end. Let τ consist of the following predicates:

$$\begin{array}{ll} U_1, U_2, P & \text{unary} \\ \sim, I, E & \text{binary.} \end{array}$$

Let φ_0 be a first-order sentence saying that

(i) the universe is partitioned into U_1, U_2, P.

(ii) $\sim$ is an equivalence relation on $U_1 \cup U_2$ such that each $\sim$-class contains exactly two elements from U_1 and U_2 each.

(iii) P and I encode the power set and elementhood over U_2 in the sense that

$$\begin{aligned}
&\forall x \forall y \big(Ixy \rightarrow (U_2 x \wedge Py)\big), \\
&\exists p \forall u \big(Pp \wedge \neg Iup\big), \\
&\forall p \forall q \big((Pp \wedge Pq \wedge \forall u(Iup \leftrightarrow Iuq)) \rightarrow p = q\big), \\
&\forall p \forall u \big(Pp \wedge U_2 u \rightarrow \exists q \big(Pq \wedge \forall w (Iwq \leftrightarrow (Iwp \vee w = u))\big)\big).
\end{aligned}$$

(iv) $E \subseteq U_1 \times U_1$ is the edge relation of a symmetric graph over U_1.

The $\{U_1, U_2, P, \sim, I\}$-reducts of the finite models of φ_0 are easily described up to isomorphism as follows. Let $U_i^{\mathfrak{A}_n} := U_i^n := \{i\} \times \{1,2\} \times n$, where n stands for the set $\{0, \ldots, n-1\}$.

Then $\mathfrak{A}_n = \big(A^n, U_1^n, U_2^n, P^n, \sim^n, I^n\big)$ where

$$\begin{aligned}
A^n &= U_1^n \cup U_2^n \cup \mathcal{P}(U_2^n), \\
P^n &= \mathcal{P}(U_2^n), \\
\sim^n &= \big\{((i,j,m),(i',j',m')) \in (U_1^n \cup U_2^n)^2 : m = m'\big\}, \\
I^n &= \big\{(u,s) \in U_2^n \times P^n : u \in s\big\}.
\end{aligned}$$

The finite models of φ_0 are of the form $\mathfrak{B} = \big(\mathfrak{A}_n, E\big)$ where E is the edge relation of a graph on U_1^n. We shall speak of $(U_1, E)^{\mathfrak{B}}$ as the graph embedded into that model $\mathfrak{B}$ of φ_0.

Lemma 6 *If ψ is any sentence of monadic second-order logic in the language $\{E\}$ of graphs, then there is a sentence ψ^* of $\mathrm{FO}[\epsilon]_{inv}$ in vocabulary τ such that the finite models of ψ^* are precisely those models of φ_0 whose embedded graph satisfies ψ.*

Proof. The following ϵ-formula $\xi(u_1, u_2)$ defines a bijection between the U_1-part and the U_2-part of any model of φ_0, and for any interpretation of ϵ:

$$\begin{aligned}\xi(u_1,u_2) = \quad & U_1u_1 \wedge U_2u_2 \wedge u_1 \sim u_2 \\ & \wedge \big(u_1 = \epsilon_x(U_1x \wedge x \sim u_1) \leftrightarrow u_2 = \epsilon_x(U_2x \wedge x \sim u_2)\big).\end{aligned}$$

Consider a formula $\chi(X)$ in a free monadic second-order variable X, which is intended to range over subsets of U_1. We want to replace X by a new first-order variable p ranging over P. Let $\chi'(p)$ be obtained from $\chi(X)$ by replacing each atom of the form Xu by the expression $\exists u_2 \xi(u, u_2) \wedge Iu_2p$. This substitution simulates X via the following parameterization of subsets of U_1: $X \subseteq U_1 \mapsto X' \subseteq U_2 \mapsto p \in P$, where X' is the image of X under the bijection defined by ξ, and p is encoding of X' according to $X' = \{u : (u,p) \in I\}$. Indeed, all finite ϵ-structures that are models of φ_0 satisfy

$$\exists X\big(X \subseteq U_1 \wedge \chi(X)\big) \longleftrightarrow \exists p\big(Pp \wedge \chi'(p)\big).$$

An inductive application of this translation to the given ψ yields an ϵ-sentence ψ' that is equivalent with the relativization of ψ to the U_1-part – for all choices of the ϵ-interpretation and in all finite models of φ_0. But then $\psi^* = \varphi_0 \wedge \psi'$ is as desired. □

Essentially the same reasoning shows that there is, for any monadic second-order definable graph query ψ, a deterministic query in FO[*witness*]$^+$ which recognizes exactly those finite models of φ_0 whose embedded graph satisfies ψ. Observe first that any finite model of φ_0 provides elements $p \in P$ encoding a subset of U_2 that contains exactly one element from each $\sim$-class over U_2, and this property of p is in first-order:

$$\eta(p) = Pp \wedge \forall u \forall u' \big((U_2u \wedge U_2u' \wedge u \sim u' \wedge u \neq u') \to \neg(Iup \leftrightarrow Iu'p)\big).$$

We apply the *witness*-operator to a formula in variables p, u_1, and u_2, where p is to be thought of as a parameter satisfying $\eta(p)$:

$$R := \mathit{witness}_{u_1} \varphi(p, u_2, u_1) \quad \text{where } \varphi(p,u_2,u_1) = Iu_2p \wedge U_1u_1 \wedge u_1 \sim u_2.$$

In this way the following formula $\xi(p, u_1, u_2)$ defines, for all p that satisfy $\eta(p)$, a bijection between U_1 and U_2:

$$\xi(p,u_1,u_2) = U_1u_1 \wedge U_2u_2 \wedge u_1 \sim u_2 \wedge (\exists u_2' Rpu_2'u_1 \leftrightarrow Iu_2p).$$

If now $\psi'(p)$ is obtained as in the proof of the last lemma, only with the parameter-dependent $\xi(p, u_1, u_2)$ in place of $\xi(u_1, u_2)$, we let $\psi^* = \varphi_0 \wedge \exists p\big(\eta(p) \wedge \psi'(p)\big)$. It is obvious that ψ^* defines a deterministic FO[*witness*]$^+$-query which simulates ψ as desired.

Corollary 7 *Any monadic second-order graph query can be simulated by a deterministic* FO[*witness*]$^+$*-query for embedded graphs in finite models of* φ_0.

We next define two infinite sequences of structures $(\mathfrak{B}_n^1)_{n\geqslant 1}$ and $(\mathfrak{B}_n^2)_{n\geqslant 1}$ such that these two sequences can be separated by a sentence of $\mathrm{FO}[\epsilon]_{\mathrm{inv}}$ but not by any sentence of plain FO. $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$ are both expansions of the above $\mathfrak{A}_n$, but with different embedded graphs (U_1, E).

$$\begin{array}{lcl} \big((1,i,m),(1,i',m)\big) \in E^{\mathfrak{B}_n^1} & \Leftrightarrow & i = i' \text{ and } m = m' \pm 1 \bmod n \\ \big((1,i,m),(1,i',m)\big) \in E^{\mathfrak{B}_n^2} & \Leftrightarrow & i = i' \text{ and } m = m' \pm 1 \ \text{ or} \\ & & i \neq i' \text{ and } \{m,m'\} = \{0, n-1\}. \end{array}$$

In $\mathfrak{B}_n^1$ the edge relation forms two disjoint cycles of length n that run in parallel through the classes of $\sim$; $\mathfrak{B}_n^2$ differs from $\mathfrak{B}_n^1$ by exactly one crossover of edges resulting in a single cycle of length $2n$.

It is clear that two cycles can be distinguished from a single cycle in monadic second-order. By the above lemma therefore there is a sentence φ in $\mathrm{FO}[\epsilon]_{\mathrm{inv}}$ such that all $\mathfrak{B}_n^1$ satisfy φ, while all $\mathfrak{B}_n^2$ satisfy $\neg\varphi$. Similarly, by the corollary, there is a deterministic $\mathrm{FO}[\mathit{witness}]^+$-query that distinguishes the $\mathfrak{B}_n^2$ from the $\mathfrak{B}_n^1$. The proof of Theorems 4 and 5 is thus reduced to the following claim.

Claim *No* FO*-sentence separates the* $\mathfrak{B}_n^1$ *from the* $\mathfrak{B}_n^2$*.*

Proof. The proof of this claim is an adaptation of the standard Ehrenfeucht-Fraïssé game arguments involving one cycle versus two cycles. For an analysis in terms of Hanf's Theorem compare e.g. [10] or [12]. Merely the fact that these two graphs appear embedded into a richer context requires a little extra care and preparation. (In particular, it should be pointed out that the Gaifman diameter of both structures is 4.)

Using an Ehrenfeucht-Fraïssé game we show that $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$ are first-order equivalent up to quantifier rank m for all sufficiently large n. We have to exhibit a strategy for the second player in the m-round game on $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$. Note that apart from the interpretation of the edge relation E over the U_1-part, $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$ are identical. Our strategy for the second player, whom we call player **II**, will be such that moves of the first player, **I**, are copied identically as long as they go outside the U_1-parts, and for moves within U_1 we let **II** always play into the same $\sim$-class as **I**. In other words, only with respect to the two members that each of these classes has in U_1 does **II** make choices.

To obtain the appropriate strategy for **II** over the U_1-parts, colour the $\sim$-classes in $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$ with n different colours $C_0, \ldots, C_{n-1}$, identically in both $\mathfrak{B}_n^i$. Considering the substructures $\mathfrak{C}_n^i = \big(\mathfrak{B}_n^i, C_0, \ldots, C_{n-1}\big) \restriction U_1$ it suffices to show for all m:

(∗) $\mathfrak{C}_n^1$ and $\mathfrak{C}_n^2$ are first-order equivalent up to quantifier rank m, for all sufficiently large n.

For then, **II** has a strategy to play the game over the U_1-part in a way that respects colours, or $\sim$-classes. A compostion of this strategy for moves in U_1 with trivial copying of moves into $U_2 \cup P$ does indeed provide a strategy in the game on on $\mathfrak{B}_n^1$ and $\mathfrak{B}_n^2$. Compatibility of the two components of this strategy follows from a simple automorphism argument. For any two choices of exactly one element from each $\sim$-class over the U_1-part — formalized by functions $\chi_1, \chi_2 \colon n \to \{1,2\}$ that encode the selection of $(1, \chi_i(k), k)$ from class $\{(1,1,k),(1,2,k)\}$ — there is an obvious automorphism ρ of $\mathfrak{A}_n$ which fixes the U_2-part and P-part point-wise and maps the one selection into the other: $\rho((1,\chi_1(k),k)) = (1,\chi_2(k),k)$ for all k. If $\overline{a}^1 \mapsto \overline{a}^2$ is a correspondence of pebbled nodes in the U_1-parts which is good for **II** in the game on the $\mathfrak{C}^i_n$, and if $\overline{c} \mapsto \overline{c}$ is the trivial correspondence produced by a copying of moves outside U_1, then $\overline{a}^1\overline{c} \mapsto \overline{a}^2\overline{c}$ is good in the game on the $\mathfrak{B}^i_n$: an automorphism of $\mathfrak{A}_n$ which takes $\overline{a}^1 \mapsto \overline{a}^2$ and fixes $\overline{c}$, shows that the correspondence $\overline{a}^1\overline{c} \mapsto \overline{a}^2\overline{c}$ respects $\tau\backslash\{E\}$, and it also respects E as $\overline{a}^1 \mapsto \overline{a}^2$ was good over the $\mathfrak{C}^i_n$.

It remains to argue for $(*)$. By Hanf's Theorem, it suffices to show that for some sufficiently large $r(m)$, ([3]), $\mathfrak{C}^1_n$ and $\mathfrak{C}^2_n$ realize exactly the same isomorphism types of Gaifman neighbourhoods of radius $r(m)$, and with the same multiplicities. The Gaifman neighbourhood of radius $r(m)$ of a fixed node b is the substructure induced on the set of nodes that can be reached from b on an $(E \cup \sim)$-path whose length is at most $r(m)$. But for $n > 2r(m)+1$ and the same fixed b, these neighbourhoods in either $\mathfrak{C}^i_n$ consist of two parallel E-strings of length $2r(m)$ (two edges short of covering the n-cycles in $\mathfrak{C}^1_n$), linked by $\sim$-edges and coloured in exactly the same way in both $\mathfrak{C}^i_n$. Therefore, the $r(m)$-neighbourhoods of the same node in $\mathfrak{C}^1_n$ and $\mathfrak{C}^2_n$, respectively, are actually isomorphic, provided only that $n > 2r(m)+1$. $\Box$

Remarks and directions for extensions

Firstly, the inseparability claim just proved can be extended to inseparability not only in FO but in the extension of FO by arbitrary monadic generalized quantifiers (or, which is the same, by arbitrary counting quantifiers). For this one can use Hella's bijective Ehrenfeucht-Fraïssé game [11] and show that the above strategy for player **II** is actually good for the (1-)bijective game on the $\mathfrak{B}^1_n$ and $\mathfrak{B}^2_n$. Here the task for **II** is harder in that, prior to each move of **I**, **II** has to propose a bijection between the universes according to which **II**'s response to **I**'s move will be made. An inspection of the strategy proposed above shows that it may be cast in this format.

Secondly, it turns out that the basic construction of the separating example can be employed for yet stronger separations if more sophisticated graph queries are implemented over the U_1-part. Consider in particular the graphs used by Cai, Fürer, and Immerman to show that graph-isomorphism is not definable in $\mathrm{C}^\omega_{\infty\omega}$, infinitary bounded-variable logic with counting quantifiers [6]. These graphs are based on gadgets whose role hinges on their behaviour under transpositions of certain 2-element subsets. Putting just those 2-element subsets into the U_1-parts of $\sim$-classes (and the rest of the graphs into a fourth part U_0 of the universe) one obtains a query in $\mathrm{FO}[\epsilon]_{\mathrm{inv}}$ that cannot even be defined in $\mathrm{C}^\omega_{\infty\omega}$, and therefore neither in fixed-point logic, nor in the counting extension of

[3] $r(m) = 2^m$ suffices according to the sharper version in [12]; $r(m) = 3^m$ is used in[10].

fixed-point logics, nor in the extension of FO by arbitrary monadic generalized quantifiers (as all of these are contained in $C^{\omega}_{\infty\omega}$).

Corollary 8 $\mathrm{FO}[\epsilon]_{inv}$ *is not even contained in* $C^{\omega}_{\infty\omega}$ *over finite structures, and therefore in particular neither in fixed-point logic (with counting) nor in the extension of first-order logic by arbitrary monadic generalized quantifiers.*

Note that these inexpressibility results carry over to $<$-invariant first-order logic, simply since $\mathrm{FO}[\epsilon]_{\mathrm{inv}} \subseteq \mathrm{FO}[<]_{\mathrm{inv}}$. We therefore obtain as a further corollary a result which has also been obtained on the basis of an inexpressibility result from [4]: $\mathrm{FO}[<]_{\mathrm{inv}}$ is not contained in the extension of FO by arbitrary monadic quantifiers, see [13]. The analogous strengthening of Theorem 5 shows that the fixed-point logics and their counting and monadic quantifier extensions do all not encompass deterministic $\mathrm{FO}[\mathit{witness}]^{+}$.

A personal note. Originally I came up with the above separating example more than six years ago, in an ill-advised attempt to separate $\mathrm{FO}[<]_{\mathrm{inv}}$ from FO — ill-advised because I was not aware of the much neater proof that Y. Gurevich had already provided for that separation. Only later did it occur to me that I had not merely reproduced a known result in a roundabout way. The issue about ϵ-invariant first-order over finite structures was first raised by X. Caicedo [7] and brought to my notice by J. Flum, compare also [8]. J. van den Bussche pointed out the connection with related issues in database theory. I would therefore like to express my thanks to these three for helping me find the right problem for my solution. My decision to submit these notes for publication, however belatedly, is due to the interest that this example has aroused on several occasions, and in various contexts, along with the versatility of the basic example that emerged partly in response to some more recent investigations and conjectures concerning both $<$-invariance and choice constructs proper, e.g. [5, 13].

References

1. S. Abiteboul, R. Hull, and V. Vianu, Foundations of Databases, Addison-Wesley, 1995.
2. S. Abiteboul, E. Simon, and V. Vianu, *Non-deterministic languages to express deterministic transformations*, in Proc. 9th ACM Symposium on Principles of Database Systems, 1990, pp. 218–229.
3. S. Abiteboul and V. Vianu, *Non-determinism in logic-based languages*, Annals of Mathematics and Artificial Intelligence, 3, 1991, pp. 151–186.
4. M. Benedikt and H. Keisler, *Expressive power of unary counters*, in: Proc. 6th International Conference on Database Theory ICDT 97, F. Afrati and P. Kolaitis (eds.), LNCS, vol. 1186, Springer-Verlag, 1997, pp. 291–305.
5. A. Blass and Y. Gurevich, *Fixed-Choice and Independent-Choice Logics*, in preparation, 1998.
6. J. Cai, M. Fürer, and N. Immerman, *An optimal lower bound on the number of variables for graph identification*, Combinatorica, 12, 1992, pp. 389–410.
7. X. Caicedo, *Hilbert's epsilon-symbol in the presence of generalized quantifiers*, Apuntes Matematicos No. 15, Universidad de los Andes 1991, also in: Quantifiers: Logics, Models and Computation, vol. 2, M. Krynicki et al. (eds.), Kluwer, 1995, pp. 469–585.
8. X. Caicedo, *Finite model theory and computational complexity*, Apuntes Matematicos No. 32, Universidad de los Andes, 1995.
9. C. Chang and H. Keisler, Model Theory, North-Holland, 1973.
10. H.-D. Ebbinghaus and J. Flum, Finite Model Theory, Perspectives in Mathematical Logic, Springer Verlag, 1996.
11. L. Hella, *Logical hierarchies in PTIME*, Information and Computation, 129, 1996, pp. 1–19.
12. N. Immerman, Descriptive Complexity, Graduate Texts in Computer Science, Springer Verlag, 1998.
13. L. Libkin and L. Wong, *Unary quantifiers, transitive closure, and relations of large degree*, Proc. Symp. Theoretical Aspects of Computer Science STACS 98, LNCS vol. 1373, Springer, 1998, pp. 183–193.