On a hierarchy involving transitive closure logic and existential second-order quantification

RICHARD L. GAULT, *Oxford University Computing Laboratory, Wolfson Building, Parks Road, Oxford OX1 3QD, UK.*
E-mail: Richard.Gault@comlab.ox.ac.uk

IAIN A. STEWART, *Department of Mathematics and Computer Science, University of Leicester, Leicester LE1 7RH, UK.*
E-mail: ias4@mcs.le.ac.uk

Abstract

We study a hierarchy of logics where each formula of each logic in the hierarchy consists of a formula of a certain fragment of transitive closure logic prefixed with an existentially quantified tuple of unary relation symbols. By playing an Ehrenfeucht-Fraïssé game specifically developed for our logics, we prove that there are problems definable in the second level of our hierarchy that are not definable in the first; and that if we are to prove that the hierarchy is proper in its entirety (or even that the third level does not collapse to the second) then we shall require substantially different constructions than those used previously to show that the hierarchy is indeed proper in the absence of the existentially quantified second-order symbols.

Keywords: Finite model theory, descriptive complexity, transitive closure logic, monadic second-order logic

1 Introduction

The problem of computing paths in directed graphs is a fundamental problem in computer science. Its logical analogue, the problem TC, *i.e.*, the class of those finite structures $\mathcal{A}$ consisting of a binary relation E and two constants C and D for which there is a path in the digraph whose edges are given by E, from the vertex C to the vertex D, has proved to be equally important in finite model theory and descriptive complexity theory. For example, extending first-order logic FO with a vectorized sequence of Lindström quantifiers corresponding to the problem TC is an elementary mechanism by which one can augment FO with a (limited) means of recursion. In what is now a seminal result, Immerman [8, 9] showed that the resulting logic, *transitive closure logic* $(\pm\mathrm{TC})^*[\mathrm{FO}]$, restricted to the class of ordered structures captures exactly the complexity class **NL** (with the corollary that **NL** is closed under complementation). Also, Fagin [5] used the problem STRCONN, *i.e.*, the class of finite structures $\mathcal{A}$ consisting of a binary relation E for which the digraph whose edges are given by E is strongly connected, to show that the monadic fragment of existential second-order logic can define problems not definable in the monadic universal fragment, and *vice versa*.

Since Immerman's and Fagin's results, transitive closure logic and (fragments of) second-order logic have subsequently been well studied. However, logics incorporating both the transitive closure operator TC and monadic second-order quantification have been barely touched upon save for Courcelle's examination [2] of the expressibility of certain "monadic second-order extensions" of transitive closure logic on graphs in relation to the logical representation of graphs (where the domain of a structure might correspond to the vertices of the graph or alternatively to the vertices and the edges).

In this paper, we study a hierarchy of logics where each formula of each logic in the hierarchy consists of a formula ψ of a certain fragment of transitive closure logic prefixed with an existentially quantified tuple of unary relation symbols (which can appear in the formula ψ). In essence, it is the hierarchy of transitive closure logic studied by Grädel in [7] (and shown there to be proper on the class of all finite structures) but with formulae prefixed with existentially quantified tuples of unary relation symbols. The question we will mainly be concerned with is: "Is the hierarchy still proper even when we allow prefixes of existentially quantified unary relation symbols?" By playing an Ehrenfeucht-Fraïssé game specifically developed for our logics, we prove that there are problems definable in the second level of our hierarchy that are not definable in the first; and that if we are to prove that the hierarchy is proper in its entirety (or even that the third level does not collapse to the second) then we shall require substantially different constructions than those currently available. Whilst our results fall short of proving that the hierarchy described above is a proper infinite hierarchy, the combinatorial complexity of our inexpressibility result testifies to the fact that playing our Ehrenfeucht-Fraïssé game is by no means straightforward (even at the lowest level of the hierarchy), and our observations as to the efficacy of our techniques for obtaining a full proper hierarchy result provide concrete evidence that new combinatorial constructions will be required for further progress to be made. Our inexpressibility result also gives rise to a proper extension of existential monadic second-order logic which is not closed under complementation.

2 Basic definitions

We begin with some definitions. The reader is referred to [3, 10] for more details. A *signature* σ is a tuple $\langle R_1, \ldots, R_r, C_1, \ldots, C_c \rangle$, where each R_i is a relation symbol, of arity $a_i > 0$, and each C_j is a constant symbol. A *finite structure* $\mathcal{A}$ *over the signature* σ, or σ*-structure*, consists of a finite *universe* or *domain* $|\mathcal{A}|$ together with a relation R_i of arity a_i, for every relation symbol R_i, and a constant $C_j \in |\mathcal{A}|$, for every constant symbol C_j (we do not generally distinguish between relations and relation symbols and between constants and constant symbols). A finite structure $\mathcal{A}$ whose domain consists of n distinct elements has *size* n, and we denote the size of $\mathcal{A}$ by $|\mathcal{A}|$ also (this does not cause confusion). We only ever consider finite structures of size at least 2, and the class of all finite structures of size at least 2 over the signature σ is denoted STRUCT(σ). A *problem* over some signature σ consists of a sub-class of STRUCT(σ) that is closed under isomorphism; that is, if $\mathcal{A}$ is in the problem then so is every isomorphic copy of $\mathcal{A}$. Throughout, all our structures are finite.

It is well-known that first-order logic, FO, is not very expressive when it comes to defining problems. One method of increasing expressibility is to extend first-order logic using a *vectorized sequence of Lindström quantifiers* corresponding to some problem Ω; or, as we prefer, an *operator* Ω for short. We shall only ever be concerned with the case when Ω is the problem TC (defined in the Introduction) over the signature σ_0 consisting of the binary relation symbol E and the two constant symbols C and D), with the resulting logic being transitive closure logic.

Transitive closure logic $(\pm\mathrm{TC})^*[\mathrm{FO}]$ consists of those formulae built using the usual constructs of first-order logic and also the operator TC, where the operator TC is applied as follows.

- Suppose that $\psi(\mathbf{x}, \mathbf{y}, \mathbf{z})$ is a formula of $(\pm\mathrm{TC})^*[\mathrm{FO}]$ such that:
  - $\mathbf{x}$ and $\mathbf{y}$ are k-tuples of distinct variables, for some $k \geq 1$;
  - $\mathbf{z}$ is an m-tuple of distinct variables, for some $m \geq 0$, each of which is different from any variable of $\mathbf{x}$ and $\mathbf{y}$; and
  - all free variables of ψ are contained in $\mathbf{x}$, $\mathbf{y}$ or $\mathbf{z}$.
- Suppose that $\mathbf{w}_0$ and $\mathbf{w}_1$ are k-tuples of variables and constant symbols (which need not be distinct).
- Then

$$\mathrm{TC}[\lambda\mathbf{x}, \mathbf{y}\psi](\mathbf{w}_0, \mathbf{w}_1)$$

  is a formula of $(\pm\mathrm{TC})^*[\mathrm{FO}]$ whose free variables are the variables of $\mathbf{z}$ together with any other variables appearing in $\mathbf{w}_0$ or $\mathbf{w}_1$.

If Φ is a sentence of the form $\mathrm{TC}[\lambda\mathbf{x}, \mathbf{y}\psi](\mathbf{w}_0, \mathbf{w}_1)$, as above, over some signature σ then we interpret Φ in a σ-structure $\mathcal{A}$ as follows (note that as Φ is a sentence, the variables of $\mathbf{z}$ are absent and the tuples $\mathbf{w}_0$ or $\mathbf{w}_1$ consist entirely of constant symbols of σ). First, we build a σ_0-structure $\Phi(\mathcal{A})$.

- The domain of the σ_0-structure $\Phi(\mathcal{A})$ is $|\mathcal{A}|^k$.
- The relation E of $\Phi(\mathcal{A})$ is defined via:
  - for any $\mathbf{u}, \mathbf{v} \in |\Phi(\mathcal{A})| = |\mathcal{A}|^k$, $E(\mathbf{u}, \mathbf{v})$ holds in $\Phi(\mathcal{A})$ if, and only if, $\psi(\mathbf{u}, \mathbf{v})$ holds in $\mathcal{A}$.
- The constants C and D of $\Phi(\mathcal{A})$ are defined via:
  - C and D are the interpretation of the tuples of constants $\mathbf{w}_0$ and $\mathbf{w}_1$, respectively, in $\mathcal{A}$.

We define that $\mathcal{A} \models \Phi$ if, and only if, $\Phi(\mathcal{A}) \in \mathrm{TC}$ (the situation where Φ has free variables is similar except that Φ is interpreted in expansions of σ-structures by an appropriate number of constants).

An alternative extension of FO is second-order logic. *Fagin's Theorem* [4] is probably the best-known result of finite model theory: a problem is in the complexity class **NP** if, and only if, it can be defined in *existential second-order logic*, Σ^1_1, *i.e.*, the formulae of the form $\exists X_1 \exists X_2 \ldots \exists X_q \varphi$, where each X_i is a relation symbol and φ is first-order. Second-order logic is very difficult to work with if one's aim is to prove inexpressibility results: however, the monadic fragment, where quantified relational symbols must be unary, has proved more amenable. *Existential monadic second-order logic*, consisting of the formulae of existential second-order logic where the quantified relation symbols are unary, is often called mon-Σ^1_1 or mon-**NP**.

Let σ be some signature and let $\mathcal{A}$ be some σ-structure. Let $X_1, X_2, \ldots, X_q$ be new unary relation symbols. We call an interpretation for each X_i over $|\mathcal{A}|$ a *colouring* of $\mathcal{A}$ (the 2^q colours involved can be considered to be the q-tuples over $\{0,1\}$ where the i^{th} bit details whether a particular domain element of $|\mathcal{A}|$ is in the relation X_i). The resulting coloured structure is the expansion of $\mathcal{A}$ by the chosen relations $X_1, X_2, \ldots, X_q$ (and so it is a $\sigma \cup \langle X_1, X_2, \ldots, X_q \rangle$-structure).

3 Our Ehrenfeucht-Fraïssé game

Our Ehrenfeucht-Fraïssé game is a natural amalgamation of the Ajtai-Fagin game from [1] and the Grädel game from [7].

DEFINITION 3.1
Let σ be some signature and let Ω be some problem over σ. The *mon-Σ_1^1-TC-game for* Ω is played between Spoiler and Duplicator, and proceeds in the following way.

- Spoiler chooses some positive integer k and a number s of sets. He then fetches k pairs of pebbles $(p_1, q_1), (p_2, q_2), \ldots, (p_k, q_k)$.
- Duplicator chooses a structure $\mathcal{A} \in \Omega$.
- Spoiler colours $\mathcal{A}$ with the s sets $X_1, X_2, \ldots, X_s$, forming the coloured structure $\mathcal{A}'$.
- Duplicator chooses a structure $\mathcal{B} \notin \Omega$.
- Duplicator colours $\mathcal{B}$ with the s sets $Y_1, Y_2, \ldots, Y_s$, and so she forms the coloured structure $\mathcal{B}'$.
- The game now proceeds in the same way as does the main part of the Grädel game from [7]. That is, there is a sequence of rounds each one of which consists of one of the following moves.
  - **A $\exists$-move.** Spoiler places a hitherto unplaced pebble p_i on an element of $|\mathcal{A}|$ and Duplicator places pebble q_i on an element of $|\mathcal{B}|$.
  - **A $\forall$-move.** Identical to a $\exists$-move but with Spoiler and Duplicator playing in the opposing structures.
  - **A TC-move.** Suppose that r pairs of pebbles have so far been placed. Spoiler chooses some $l \leq (k-r)/2$ and selects a sequence $\mathbf{u_0}$, $\mathbf{u}_1$, ..., $\mathbf{u}_m$ of l-tuples over $|\mathcal{A}|$, for some $m \geq 1$, such that $\mathbf{u}_0$ and $\mathbf{u}_m$ consist entirely of constants and previously pebbled elements. Duplicator then selects a sequence $\mathbf{v}_0$, $\mathbf{v}_1$, ..., $\mathbf{v}_n$ of l-tuples over $|\mathcal{B}|$, for some $n \geq 1$, such that $\mathbf{v}_0$ and $\mathbf{v}_n$ are the analogous tuples over $|\mathcal{B}|$ to $\mathbf{u}_0$ and $\mathbf{u}_m$ (that is, given by the corresponding constants and pebbles). Spoiler then places $2l$ hitherto unplaced q-pebbles on the elements of the tuples $\mathbf{v}_i$ and $\mathbf{v}_{i+1}$, for some $i \in \{0, 1, \ldots, n-1\}$ (of course, some pebbles might end up being placed on the same domain element); and Duplicator replies by placing the $2l$ corresponding p-pebbles on the elements of the tuples $\mathbf{u}_j$ and $\mathbf{u}_{j+1}$, for some $j \in \{0, 1, \ldots, m-1\}$.
  - **A $\neg$TC-move.** Identical to a TC-move but with Spoiler and Duplicator playing in the opposing structures.
- The game continues until all $2k$ pebbles have been placed.

Duplicator wins a play of the game if, and only if, at the end of the play the mapping defined by the pebbles and the constants of the coloured structures is a (well-defined) partial isomorphism from $\mathcal{A}'$ to $\mathcal{B}'$ (where this partial isomorphism must respect colour). ∎

The traditional Ehrenfeucht-Fraïssé game is a restriction of that in Definition 3.1 in that there is no colouring phase and the only moves allowed are $\exists$-moves and $\forall$-moves. Our game has been designed for the following logics; which are extensions of fragments of transitive-closure logic by existential second-order quantification.

DEFINITION 3.2
For every string ω over the alphabet $\{\exists, \forall, \mathrm{TC}, \neg\mathrm{TC}\}$, we define the quantifier class $\mathrm{TC}^0(\omega)$ in $(\pm\mathrm{TC})^*[\mathrm{FO}]$ inductively as follows.

- The class $\mathrm{TC}^0(\epsilon)$ contains the quantifier-free first-order formulae.
- For a quantifier $Q \in \{\exists, \forall\}$, the class $\mathrm{TC}^0(Q\omega)$ is the closure under conjunctions and disjunctions of the class of formulae $\mathrm{TC}^0(\omega) \cup \{(Qx)\varphi : x \text{ is some variable and } \varphi \in \mathrm{TC}^0(\omega)\}$.
- The class $\mathrm{TC}^0(\mathrm{TC}\omega)$ is the closure under conjunctions and disjunctions of the class of formulae of $\mathrm{TC}^0(\omega)$ and formulae of the form $\mathrm{TC}[\lambda\mathbf{x}, \mathbf{y}\varphi](\mathbf{z}, \mathbf{w})$, where $\varphi \in \mathrm{TC}^0(\omega)$.
- The class $\mathrm{TC}^0(\neg\mathrm{TC}\omega)$ is the closure under conjunctions and disjunctions of the class of formulae of $\mathrm{TC}^0(\omega)$ and formulae of the form $\neg\mathrm{TC}[\lambda\mathbf{x}, \mathbf{y}\varphi](\mathbf{z}, \mathbf{w})$, where $\varphi \in \mathrm{TC}^0(\omega)$. ∎

As remarked in [7], any formula of $(\pm\mathrm{TC})^*[\mathrm{FO}]$ is equivalent to a formula in some class $\mathrm{TC}^0(\omega)$, where $\omega \in \{\mathrm{TC}, \neg\mathrm{TC}\}^*$.

DEFINITION 3.3
For every string ω over the alphabet $\{\exists, \forall, \mathrm{TC}, \neg\mathrm{TC}\}$ and for every $q \geq 1$, we define the quantifier class mon-Σ^1_1-$\mathrm{TC}^q(\omega)$ as

$$\{\exists X_1 \exists X_2 \ldots \exists X_q \varphi : \text{ each } X_i \text{ is a unary relation symbol and } \varphi \in \mathrm{TC}^0(\omega)\},$$

with mon-Σ^1_1-$\mathrm{TC}^0(\omega)$ defined as $\mathrm{TC}^0(\omega)$. We set

$$\text{mon-}\Sigma^1_1\text{-TC}(\omega) = \cup_{q=1}^{\infty} \text{mon-}\Sigma^1_1\text{-TC}^q(\omega);$$

and

$$\text{mon-}\Sigma^1_1\text{-TC} = \{\psi : \psi \in \text{mon-}\Sigma^1_1\text{-TC}(\omega), \text{ for some } \omega \in \{\exists, \forall, \mathrm{TC}, \neg\mathrm{TC}\}^*\}.$$

∎

Our game, above, gives rise to the following proposition, whose proof follows from that of Theorem 3.5 below.

PROPOSITION 3.4
Duplicator has a winning strategy for the mon-Σ^1_1-TC-game for Ω if, and only if, Ω is not definable in mon-Σ^1_1-TC. ∎

Proposition 3.4 is just a special case of a more general result. For any word $\omega \in \{\forall, \exists, \mathrm{TC}, \neg\mathrm{TC}\}^*$ and for any $q \geq 0$, we may define the mon-Σ^1_1-$\mathrm{TC}^q(\omega)$-game to be the same as the mon-Σ^1_1-TC-game, save that Spoiler is obliged to choose q as his value of s and to choose his moves according to the string ω, *i.e.*, the i^{th} move of the play is determined by the i^{th} symbol of ω; and so the string ω also dictates the number of rounds in any play. In any play of the game, it is always the case that all pebbles are placed (and so it may be the case that not every choice for k is legitimate: it all depends on ω). The mon-Σ_1^1-TC(ω)-game is defined as expected (where Spoiler can choose any value for s). Then we have the following result.

THEOREM 3.5
Let Ω be a problem over some signature σ. Let $\omega \in \{\forall, \exists, TC, \neg TC\}^*$ and let $q \geq 0$. Then Duplicator has a winning strategy in the mon-Σ_1^1-$TC^q(\omega)$-game (resp. mon-Σ_1^1-$TC(\omega)$-game) for Ω if, and only if, Ω is not definable in mon-Σ_1^1-$TC^q(\omega)$ (resp. mon-Σ_1^1-$TC(\omega)$).

PROOF. We shall only prove the theorem in the case where q is given. The case where q is unrestricted is similar.

We begin by proving the "if" direction. This is the most useful direction in practice and is the easier to prove. In fact, we prove the contrapositive: that is, if Ω is definable by a sentence of mon-Σ_1^1-TC$^q(\omega)$ then Duplicator does not have a winning strategy.

Suppose that Ω is defined by the sentence $\psi = \exists X_1 \exists X_2 \ldots \exists X_q \varphi$, where $\varphi \in$ TC$^q(\omega)$. During the first step of the game, Spoiler is obliged to take $s = q$. Duplicator responds by choosing some structure $\mathcal{A} \in \Omega$ which Spoiler must colour. As $\mathcal{A} \models \psi$, there must be some assignment to the sets $X_1, X_2, \ldots, X_q$ so that the colouring $\mathcal{A}'$ of $\mathcal{A}$ by this assignment satisfies $\mathcal{A}' \models \varphi$. Spoiler chooses this colouring to obtain the coloured structure $\mathcal{A}'$.

Duplicator now responds by choosing some $\mathcal{B} \notin \Omega$ and colouring it. Since $\mathcal{B} \not\models \psi$, no matter how she colours $\mathcal{B}$, to obtain the coloured structure $\mathcal{B}'$, it is always the case that $\mathcal{B}' \not\models \varphi$.

The game now proceeds as does the Grädel game on Ω. The proof that Spoiler has a winning strategy (and hence that Duplicator does not) may be found in [7]. We need merely observe that the proof presented in [7] still goes through when the structures in question are coloured.

For the converse, we will prove that if Duplicator does not have a winning strategy in the mon-Σ_1^1-TC$^q(\omega)$-game for Ω then Ω is definable in the corresponding logic.

If Duplicator does not have a winning strategy then Spoiler must have such a strategy. That is, whichever $\mathcal{A} \in \Omega$ Duplicator chooses, it may be coloured by the q sets in such a way that whichever $\mathcal{B} \notin \Omega$ Duplicator chooses, and however she colours it, Spoiler has a winning strategy in the Grädel-part of the game. For any $\mathcal{A}$, denote Spoiler's winning colouring of $\mathcal{A}$ by $\mathcal{A}'$.

Since the proof of correctness of the Grädel game goes through even when the structures are coloured, it follows that whichever $\mathcal{A} \in \Omega$ Duplicator chooses as her first structure, and whichever coloured $\mathcal{B}'$ she chooses as her second structure, there is at least one formula $\varphi_{\mathcal{A},\mathcal{B}'} \in$ TC$^q(\omega)$ such that $\mathcal{A}' \models \varphi_{\mathcal{A},\mathcal{B}'}$ and $\mathcal{B}' \models \neg\varphi_{\mathcal{A},\mathcal{B}'}$. Of course, $\varphi_{\mathcal{A},\mathcal{B}'}$ will make use of up to q free monadic second-order variables. Let $\Phi_{\mathcal{A}}$ be the conjunction of these $\varphi_{\mathcal{A},\mathcal{B}'}$ over all possible choices of $\mathcal{B}'$. This conjunction is permitted since there are only a finite number of such formulae up to logical equivalence. Note that $\Phi_{\mathcal{A}} \in$ TC$^q(\omega)$.

Now, clearly $\mathcal{A}' \models \Phi_{\mathcal{A}}$. Also, for every $\mathcal{B} \notin \Omega$ and for every possible colouring $\mathcal{B}'$ of $\mathcal{B}$, $\mathcal{B}' \models \neg\Phi_{\mathcal{A}}$. So let Ψ be the disjunction of the formulae of $\{\Phi_{\mathcal{A}} : \mathcal{A} \in \Omega\}$. Once again, this is permitted since there are only a finite number of such formulae up to logical equivalence. Then $\Psi \in \mathrm{TC}^q(\omega)$ and, moreover, for any $\mathcal{A} \in \Omega$, $\mathcal{A}' \models \Psi$. Furthermore, for all $\mathcal{B} \notin \Omega$ and for any colouring $\mathcal{B}'$ of $\mathcal{B}$, $\mathcal{B}' \models \neg\Psi$.

So for every $\mathcal{A} \in \Omega$:

$$\mathcal{A} \models \exists X_1 \exists X_2 \ldots \exists X_q \Psi$$

where the X_i are chosen to be the free second-order variables of Ψ. Similarly, for every $\mathcal{B} \notin \Omega$:

$$\mathcal{B} \models \neg\exists X_1 \exists X_2 \ldots \exists X_q \Psi.$$

Hence, $\exists X_1 \exists X_2 \ldots \exists X_q \Psi$ is a mon-Σ^1_1-$\mathrm{TC}^q(\omega)$ formula which defines Ω. ∎

4 Playing our game

In [8], Immerman exhibited a particularly strong normal form result for formulae of transitive closure logic on the class of ordered structures. This normal form result was established by induction on the symbolic complexity of a formula. In fact, his case-by-case analysis goes through even on unordered structures except for the elimination of the universal quantifier. With this observation in mind, consider the following hierarchy.

Definition 4.1
The hierarchy

$$\mathrm{TC}(0) \subseteq \forall\text{-}\mathrm{TC}(0) \subseteq \mathrm{TC}(1) \subseteq \forall\text{-}\mathrm{TC}(1) \subseteq \mathrm{TC}(2) \subseteq \ldots$$

within transitive closure logic is defined as follows.

- $\mathrm{TC}(0)$ consists of all formulae of the form $\mathrm{TC}[\lambda \mathbf{x}, \mathbf{y}\varphi](\mathbf{w}_0, \mathbf{w}_1)$, where φ is first-order (and may contain free variables other than those of $\mathbf{x}$ and $\mathbf{y}$).
- $\forall$-$\mathrm{TC}(m)$ is the universal closure of $\mathrm{TC}(m)$, *i.e.*, the set of formulae of the form $\forall x_1 \forall x_2 \ldots \forall x_m \varphi$, where $\varphi \in \mathrm{TC}(m)$.
- $\mathrm{TC}(m+1)$ is the set of formulae of the form $\mathrm{TC}[\lambda \mathbf{x}, \mathbf{y}\varphi](\mathbf{w}_0, \mathbf{w}_1)$, where $\varphi \in \forall$-$\mathrm{TC}(m)$. ∎

Grädel proved the following result [7].

Theorem 4.2
Over signatures containing at least two constant symbols which are always interpreted differently in any structure, $TC(m) \subset \forall$-$TC(m) \subset TC(m+1)$, for all $m \geq 0$. ∎

One can ask whether Grädel's hierarchy result holds if one prefixes the formulae of the logics $\mathrm{TC}(i)$ and $\forall$-$\mathrm{TC}(i)$ with existentially quantified prefixes of (new) relation symbols.

Definition 4.3
Define

$$\text{mon-}\Sigma^1_1\text{-}\mathrm{TC}(i) = \{\exists X_1 \exists X_2 \ldots \exists X_q \varphi : q \geq 0, \text{ each } X_j \text{ is a unary relation symbol and } \varphi \in \mathrm{TC}(i)\};$$

and

$$\text{mon-}\Sigma_1^1\text{-}\forall\text{-TC}(i) \quad = \quad \{\exists X_1 \exists X_2 \ldots \exists X_q \varphi : q \geq 0, \text{ each } X_j \text{ is a unary relation symbol and } \varphi \in \forall\text{-TC}(i)\}.$$

The resulting hierarchy is

$$\text{mon-}\Sigma_1^1\text{-TC}(0) \subseteq \text{mon-}\Sigma_1^1\text{-}\forall\text{-TC}(0) \subseteq \text{mon-}\Sigma_1^1\text{-TC}(1) \subseteq \text{mon-}\Sigma_1^1\text{-}\forall\text{-TC}(1) \subseteq \ldots$$

∎

Note that mon-Σ_1^1-TC(0) is a strict extension of mon-Σ_1^1 (TC can be defined in the former logic but not the latter [1]).

Let us now make a start on answering the question posed prior to Definition 4.3. In order to make developing winning strategies in our games easier, we shall use a result due to Fagin, Stockmeyer and Vardi [6].

For any structure $\mathcal{A}$, we say that $a, b \in |\mathcal{A}|$ are *adjacent* if either $a = b$ or there is some tuple $\mathbf{t}$ in some relation of $\mathcal{A}$ such that both a and b appear as elements of $\mathbf{t}$. The *degree* of an element $a \in |\mathcal{A}|$ is the number of elements of $|\mathcal{A}|$ adjacent to a but not equal to a. For any integer $d \geq 1$ and $a \in |\mathcal{A}|$, we define the *neighbourhood of radius d about a*, $\text{Nbd}(d, a)$, recursively as follows.

$$\begin{aligned} \text{Nbd}(1, a) &= \{a\}. \\ \text{Nbd}(d+1, a) &= \{b \in |\mathcal{A}| : b \text{ is adjacent to some } b' \in \text{Nbd}(d, a)\}. \end{aligned}$$

The *d-type* of an element $a \in |\mathcal{A}|$ is the isomorphism type of $\text{Nbd}(d, a)$ (where a is regarded as a distinguished constant).

Crucial to Fagin, Stockmeyer and Vardi's result, alluded to above, is the notion of (d, m)-equivalence. Let d and m be non-zero natural numbers. Two structures $\mathcal{A}$ and $\mathcal{B}$ over the same signature are *(d, m)-equivalent* if for any d-type τ, either $\mathcal{A}$ and $\mathcal{B}$ have exactly the same number of elements of d-type τ or they both have at least m elements of d-type τ. Fagin, Stockmeyer and Vardi proved the following [6].

THEOREM 4.4
Let σ be some signature containing just relation symbols (and no constant symbols). Let l and f be non-zero natural numbers. There there exist non-zero natural numbers d and m where d depends solely on l and m depends only on l and f such that whenever $\mathcal{A}$ and $\mathcal{B}$ are (d, m)-equivalent and every element of both structures has degree at most f then Duplicator has a winning strategy in the traditional Ehrenfeucht-Fraïssé game in which Spoiler is obliged to choose l as his value of k. ∎

Theorem 4.4 can often be used to simplify the proof that Duplicator has a winning strategy in an Ehrenfeucht-Fraïssé game. For instead of actually having to exhibit a strategy, it is sufficient to prove that the two structures chosen by Duplicator are, in some sense, "locally equivalent". This is often a considerably easier task.

In fact, Theorem 4.4 also holds over signatures which contain constant symbols (as remarked in [11]). To see this, let $\sigma = \langle R_1, R_2, \ldots, R_r, C_1, C_2, \ldots, C_c \rangle$ be a signature containing the c constant symbols $C_1, C_2, \ldots C_c$, and let $\mathcal{A}$ be a σ-structure. Consider now the purely relational signature $\sigma' = \langle R_1, R_2, \ldots, R_r, R'_1, R'_2, \ldots, R'_c \rangle$ in which each constant symbol of σ has been replaced by a unary relation symbol. We define the structure $\mathcal{A}'$ over σ' in the following way. The universe of $\mathcal{A}'$ is equal to the universe of $\mathcal{A}$. Each relation R_i of $\mathcal{A}'$ is identical to the relation R_i of $\mathcal{A}$. Each relation R'_i of $\mathcal{A}'$ is defined so that $R'_i(u)$ holds for precisely one value $u \in |\mathcal{A}'|$; specifically, for that value for which $(\mathcal{A}, u) \models C_i = x$. Clearly, for any d and m, a pair of structures $\mathcal{A}$ and $\mathcal{B}$ are (d, m)-equivalent if, and only if, the corresponding structures $\mathcal{A}'$ and $\mathcal{B}'$ are (d, m)-equivalent. Furthermore, Duplicator has a winning strategy in the traditional Ehrenfeucht-Fraïssé game on $\mathcal{A}$ and $\mathcal{B}$ if, and only if, she has a winning strategy on $\mathcal{A}'$ and $\mathcal{B}'$.

THEOREM 4.5
Let $\sigma = \langle E, C, D\rangle$, where E is a binary relation symbol and C and D are constant symbols. Let Ω be the problem consisting of those σ-structures which when considered as undirected graphs, in the natural way, are connected. Then Ω cannot be defined by a sentence of the logic mon-Σ^1_1-$TC(0)$.

PROOF. We shall play a mon-Σ^1_1-TC(ω)-game where ω is a string over the alphabet $\{\forall, \exists, \mathrm{TC}, \neg\mathrm{TC}\}$ of the form $\mathrm{TC}\omega'$, with $\omega' \in \{\forall, \exists\}^*$. Spoiler begins by choosing positive integers s and k. Of course, the value of k will determine the length of tuples chosen by Spoiler in the first TC-move of the Grädel-part of any play of the game: call this number k_0 and let $k_1 = k - k_0$ (and so k_1 is the length of ω'). Duplicator takes $f = 2$ and $l = k_1$ and applies Theorem 4.4 to obtain d and m. Duplicator begins by choosing the structure $\mathcal{A}$ to be a cycle of a certain length p, together with two constants C and D (the structure $\mathcal{B}$ will ultimately be a pair of disjoint cycles, hence our choice of $f = 2$). The choice of the actual value for p is made as follows.

Given any undirected path of length $2d$ (with $2d + 1$ vertices), there are exactly $(2^s)^{2d+1}$ ways of colouring it (with 2^s colours). Consequently, for any r, if Duplicator chooses $p \geq r(2d + 1)2^{s(2d+1)}$ then no matter how Spoiler might colour the cycle $\mathcal{A}$, the resulting coloured cycle $\mathcal{A}'$ must contain at least r identically coloured, non-overlapping regions of length $2d$ (any region includes $2d + 1$ vertices, so there are at least $r2^{s(2d+1)}$ batches of $2d + 1$ vertices on the cycle; and thus some colouring of a region appears at least r times). Call these regions $\Delta_1, \Delta_2, \ldots, \Delta_{r'}$, with $r' \geq r$. Note that in fact p may be chosen to be much smaller than this and this property would still hold: we are interested only in the existence of such a p, not in a minimal value.

We must now choose a suitable value for r. The number of distinct isomorphism types of radius d on the cycle $\mathcal{A}$ (ignoring, for the moment, the constants C and D which must also eventually be chosen) is at most $2^{s(2d+1)}$. So the power set of such isomorphism types has size at most $2^{2^{s(2d+1)}}$. Consider any "gap" between two regions Δ_i and Δ_{i+1}, where $i \in \{1, 2, \ldots, r' - 1\}$, or between $\Delta_{r'}$ and Δ_1 (note that no vertex lies in both a gap and a region). Every gap gives rise to the subset of isomorphism types of radius d, i.e., d-types, witnessed by the vertices of the gap. If we choose $r \geq 2(m + 2k_0 + 2)2^{2^{s(2d+1)}}$ then there exist at least $2(m + 2k_0 + 2)$ gaps giving rise to the same set S of isomorphism types of radius d. Moreover, any one of the isomorphism types in S is witnessed by at least $2(m + 2k_0 + 2)$ vertices in $\mathcal{A}'$. So, Duplicator chooses $p \geq 2(m + 2k_0 + 2)2^{2^{s(2d+1)}}(2d + 1)2^{s(2d+1)}$.

Consider the following strategy of Duplicator. Duplicator chooses a cycle $\mathcal{A}$ of length p, together with two constants C and D, and Spoiler replies by colouring this cycle, with 2^s colours, to give the coloured cycle $\mathcal{A}'$. Duplicator takes a copy of $\mathcal{A}'$

and determines the r' identically coloured regions and (at least) $2(m + 2k_0 + 2)$ gaps giving rise to the same set of isomorphism types of radius d. Duplicator chooses a gap, without loss of generality between the regions Δ_1 and Δ_2, say, such that neither of the constants C and D lie in the gap nor in the adjacent Δ-regions. Then Duplicator copies the portion of the coloured cycle from the median vertex of the region Δ_i through the gap to the median vertex of the region Δ_{i+1}; and then Duplicator fuses the two extreme vertices of this path (that is, the copies of the median vertices) to form a coloured cycle. This process can be visualised in Fig. 1. Duplicator chooses the coloured structure $\mathcal{B}'$ to be the disjoint union of the copy of the coloured cycle $\mathcal{A}'$ (complete with constants) and the newly-formed coloured cycle. Denote by $\mathcal{B}'_0$ the coloured cycle of $\mathcal{B}'$ that is isomorphic to $\mathcal{A}'$; and denote by $\mathcal{B}'_1$ the other coloured cycle.

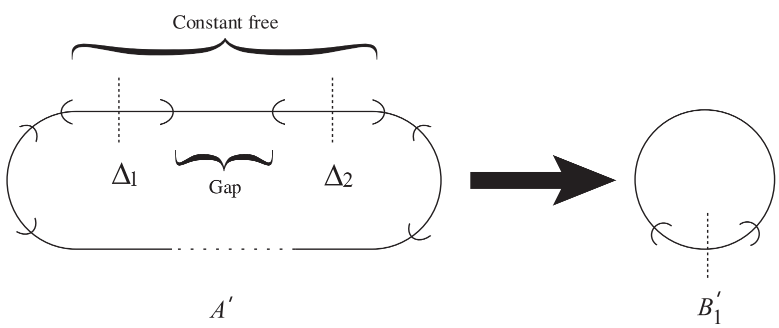

Figure 1. Forming the new cycle.

According to the game, Spoiler chooses a path of k_0-tuples over $|\mathcal{A}'|$. Duplicator's strategy is such that she replies with the isomorphic path of k_0-tuples over $|\mathcal{B}'_0|$. Whichever pair of tuples Spoiler places pebbles on, Duplicator responds isomorphically. They now play a traditional Ehrenfeucht-Fraïssé game, determined by ω, on the structures $\mathcal{A}'$ and $\mathcal{B}'$ where each has been augmented with $2k_0$ constants. Let us consider the isomorphism types of radius d in each structure. Obviously, the isomorphism types arising from $\mathcal{A}'$ and $\mathcal{B}'_0$ are identical; and there are exactly the same numbers of vertices of each present. As regards any isomorphism type arising from a vertex of $\mathcal{B}'_1$, there are at least m vertices of identical type in $\mathcal{A}$. To see this, note that ignoring the constants C and D and the $2k_0$ pebbles, there are at least $m + 2k_0 + 2$ vertex-disjoint portions of $\mathcal{A}'$ of the form $(\Delta_i, \text{gap}, \Delta_{i+1})$ giving rise to the same set of isomorphism types as witnessed by the vertices of $|\mathcal{B}'_1|$. Taking into account C and D and the $2k_0$ pebbles results in at least m such portions. Hence, $\mathcal{A}'$ and $\mathcal{B}'$ are (d, m)-equivalent and Duplicator has a winning strategy by Theorem 4.4 (note that we are applying Theorem 4.4 to coloured structures). Theorem 3.5 implies that Ω is not definable in mon-Σ^1_1-TC(0) and the result follows. ∎

Note that we deliberately chose our signature in Theorem 4.5 to contain (distinct) constant symbols so that we could actually define sentences of mon-Σ-1-TC(0).

Corollary 4.6
mon-Σ-1-TC(0) $\subset$ mon-Σ-1-$\forall$-TC(0); and mon-Σ-1-TC(0) is not closed under complementation.

PROOF. Follows immediately from Theorem 4.5 and the fact [5] that the problem consisting of those undirected graphs that are not connected is definable in mon-Σ^1_1; and so in mon-Σ-1-TC(0). ∎

Of course, we would like to obtain that the hierarchy in Definition 4.3 is proper. However, Grädel's constructions from [7], to prove his hierarchy result in the absence of second-order quantification, cannot be used as we now explain. We have already commented (prior to Definition 3.3) that existential first-order quantification can be simulated using a TC quantifier. One consequence of this is that in order to prove that $\forall$-TC$(m) \subseteq$ TC$(m+1)$, Grädel merely needed to show that $\forall$-TC$(m) \subseteq \exists$-$\forall$-TC(m) (the definition of this latter logic should be clear). However, when we have a second-order $\exists$ quantifier present, this is not the case. For a sentence of mon-Σ^1_1-$\exists$-$\forall$-TC(m) of the form

$$\exists X_1 \exists X_2 \ldots \exists X_q \exists y \forall z_1 \forall z_2 \ldots \forall z_m \varphi,$$

with $\varphi \in$ TC(m), is equivalent to one of the form

$$\exists X_1 \exists X_2 \ldots \exists X_q \exists Y (\text{“}Y \neq \emptyset\text{”} \wedge \forall y(\neg Y(y) \vee \forall z_1 \forall z_2 \ldots \forall z_m \varphi)).$$

At first sight, it may appear that we have not gained very much: we still require a first-order existential quantifier to express “$Y \neq \emptyset$”. Note, however, that “$Y \neq \emptyset$” $\in$ TC(0), and consequently may be moved inside φ. Thus our original sentence is equivalent to one from mon-Σ^1_1-$\forall$-TC(m). As a consequence, any attempt to prove a full hierarchy result analogous to Grädel's must make more essential use of the quantifier TC than its ability to define first-order existential quantification. This appears to be a combinatorially difficult problem to circumvent.

As a final remark, note that in our proof of Theorem 4.5, we make no real use of the quantifier TC beyond the fact that the problem TC is *closed under superstructures*; that is, if $\mathcal{A}$ and $\mathcal{B}$ are structures over the signature $\langle E, C, D\rangle$ such that $|\mathcal{A}| \subseteq |\mathcal{B}|$, $\mathcal{A}$ is $\mathcal{B}$ restricted to $|\mathcal{A}|$ and $\mathcal{A} \in$ TC then necessarily $\mathcal{B} \in$ TC. If we were to define a logic mon-Σ^1_1-Γ, for some problem Γ that is closed under superstructures, just as we have defined mon-Σ^1_1-TC (complete with associated fragments and analogous semantics), the proof of Theorem 4.5 essentially goes through so that we obtain the following more general result.

COROLLARY 4.7
Let $\sigma = \langle E, C, D\rangle$, where E is a binary relation symbol and C and D are constant symbols. Let Ω be the problem consisting of those σ-structures which when considered as undirected graphs, in the natural way, are connected. Let Γ be some problem that is closed under superstructures. Then Ω cannot be defined by a sentence of the logic mon-Σ^1_1-$\Gamma(0)$. ∎

References

[1] M. Ajtai and R. Fagin, Reachability is harder for directed than for undirected graphs, *Journal of Symbolic Logic* **55** (1990) 1–48.

[2] B. Courcelle, On the expression of graph properties in some fragments of monadic second-order logic, *DIMACS Series in Discrete Mathematics and Theoretical Computer Science* **31** (1997) 33–62.

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

[4] R. Fagin, Generalized first-order spectra and polynomial-time recognizable sets, in: *Complexity of Computation* (ed. R.M. Karp), Volume 7 of SIAM-AMS Proceedings (1974) 43–73.

[5] R. Fagin, Monadic generalized spectra, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **21** (1975) 89–96.

[6] R. Fagin, L. Stockmeyer and M.Y. Vardi, On monadic NP vs. monadic co-NP, *Information and Computation* **120** (1995) 78–92.

[7] E. Grädel, On transitive closure logic, *Proceedings of Computer Science Logic CSL 93* (eds. E. Börger, Y. Gurevich and K. Meinke), *Lecture Notes in Computer Science* 626, Springer-Verlag (1991) 149–163.

[8] N. Immerman, Languages that capture complexity classes, *SIAM Journal of Computing* **16** (1987) 760–778.

[9] N. Immerman, Nondeterministic space is closed under complementation, *SIAM Journal of Computing* **17** (1988) 935–938.

[10] N. Immerman, *Descriptive Complexity*, Springer-Verlag (1998).

[11] T. Schwentick, On winning Ehrenfeucht games and monadic NP, *Annals of Pure and Applied Logic* **79** (1996) 61–92.

Received 30 August 2001.