ELSEVIER

Theoretical Computer Science 174 (1997) 97–121

Theoretical Computer Science

On winning strategies in Ehrenfeucht–Fraïssé games

Sanjeev Arora [a,1], Ronald Fagin [b,*]

[a] *35 Olden Street, Princeton NJ 08544, USA*
[b] *650 Harry Road, San Jose, California 95120-6099, USA*

Received May 1994; revised December 1995
Communicated by P. Kanellakis

Abstract

We present a powerful and versatile new sufficient condition for the second player (the "duplicator") to have a winning strategy in an Ehrenfeucht–Fraïssé game on graphs. We accomplish two things with this technique. First, we give a simpler and much easier-to-understand proof of Ajtai and Fagin's result that reachability in directed finite graphs is not in monadic NP. (Monadic NP, otherwise known as monadic Σ_1^1, corresponds to existential second-order logic with the restriction that the second-order quantifiers range only over sets, and not over relations of higher arity, such as binary relations.) Second, we show that this result holds in the presence of a larger class of built-in relations than was known before.

1. Introduction

The *computational complexity* of a problem is the amount of resources, such as time or space, required by a machine that solves the problem. The *descriptive complexity* of a problem is the complexity of describing the problem in some logical formalism [18]. The two complexities are sometimes related. This was first discovered by Fagin, who showed [9] that the complexity class NP coincides with the class of properties of finite structures expressible in existential second-order logic, otherwise known as Σ_1^1. Consequently, NP$=$co-NP if and only if existential and universal second-order logic have the same expressive power over finite structures, i.e., if and only if $\Sigma_1^1 = \Pi_1^1$. In a similar vein, Immerman and Vardi proved that the complexity class P coincides with the class of properties of finite ordered structures expressible in fixpoint logic [17, 22].

The famous conjectures of computational complexity theory, such as NP $\neq$ co-NP, seem difficult to prove. Therefore their connection to questions in descriptive complexity holds the promise that techniques from one domain can be brought to bear on

* Corresponding author. email: fagin@almaden.ibm.com.
[1] This work was done while the author was at UC Berkeley and visiting IBM Almaden. email: arora@cs.princeton.edu.

PII S0304-3975(96)00015-1

questions in the other domain. In particular, there is a standard technique in finite-model theory for proving separation results: Ehrenfeucht–Fraïssé games (see Section 3). Fagin showed that $\Sigma_1^1 \neq \Pi_1^1$ if and only if such a separation can be proven via second-order Ehrenfeucht–Fraïssé games [10]. In the same paper, Fagin also suggested that partial progress on the $\Sigma_1^1 \neq \Pi_1^1$ question could be made by restricting the expressive power of these classes: instead of allowing second-order quantification (i.e., over relations of arbitrary arity), allow quantification only over sets. Since quantifying over sets corresponds to quantification over *monadic* relations, Fagin et al. [12] term these restricted classes as *monadic NP* and *monadic co-NP* respectively, a terminology that has since gained acceptance. Note that, in spite of its seemingly restrictive syntax, monadic NP contains nontrivial problems, including NP-complete ones such as 3-colorability and satisfiability.

Fagin [10] showed that monadic NP $\neq$ monadic co-NP. Specifically, he showed that connectivity of finite graphs – a property that is easily seen to be in monadic co-NP – is not in monadic NP. His proof uses a certain Ehrenfeucht–Fraïssé game on graphs. Such games have since been used in many other nonexpressibility results.

An *Ehrenfeucht–Fraïssé game* is played between two players, called the spoiler and the duplicator; nonexpressibility results involve proving that the duplicator has a winning strategy. Often this proof of existence is quite complicated (see e.g. [10, 1]), which is not surprising: a spoiler's strategy may be arbitrarily complicated, and the proof has to argue that *no* strategy prevails against the duplicator. Thus, it seems quite important to develop tools for proving that the duplicator wins.

Several such conditions have been identified. Among these are (a) a formulation of second-order Ehrenfeucht–Fraïssé games by Ajtai and Fagin [1], for which it seems easier to prove that the duplicator has a winning strategy; (b) a sufficient condition (due essentially to Hanf [16]) for the duplicator to have a winning strategy; and (c) the idea of playing Ehrenfeucht–Fraïssé games over random structures. Techniques (a) and (c) were used by Ajtai and Fagin [1], and all three techniques were used by Fagin et al. [12].

Thus a "library" of tools seems to be emerging, each of which further simplifies the task of showing that the duplicator has a winning strategy. Clearly, it is important to enlarge this library with tools that are as intuitive and natural as possible, so that researchers have an easier time in proving nonexpressibility results. To give an example, the three "library" tools mentioned above were used by Fagin et al. as follows: (1) they give a simple proof (much simpler than Fagin's) that connectivity is not in monadic NP, and (2) they show that connectivity is not in monadic NP in the presence of a large class of built-in relations (this was previously known only for a built-in successor relation [6]; we discuss built-in relations shortly).

In this paper we give a new sufficient condition for the duplicator to have a winning strategy. We hope that this condition will prove useful and intuitive. We use it as follows. (1) We give a proof that directed reachability (given a directed graph and two distinguished vertices s and t in it, the problem of deciding whether there is a directed path from s to t) is not in monadic NP. The proof is much easier than the earlier proof by Ajtai and Fagin [1]. (2) We show that directed reachability is not in monadic NP in the presence of a larger class of built-in relations than was known before. (3) We note that the condition can also be used in place of Hanf's condition in proof of Fagin et al. for the connectivity problem. (It is arguable, though, whether this makes the proof simpler.)

There are several reasons for interest in connectivity and directed reachability, the two problems considered in this paper. Cosmadakis [5] has shown that connectivity reduces (via a very weak kind of reduction) to a host of other problems, including non-3-colorability. So the fact that connectivity is not in monadic NP implies that these problems are not in monadic NP. Thus, connectivity seems to have a special significance.

Directed reachability is an interesting problem because it is known in some senses to be a more difficult problem to deal with than connectivity. For instance, a surprising result of Kanellakis (private communication, 1986; see also [2]) says that *undirected* reachability (where we consider only undirected graphs) is in monadic NP. (In contrast, undirected connectivity is not in monadic NP, as mentioned above.) Furthermore, Cosmadakis has shown that connectivity does not reduce, in his sense, to directed reachability. This suggests that if directed reachability is not in monadic NP, then proving this requires techniques different from those used to show that connectivity is not in monadic NP. Ajtai and Fagin [1] resolved the status of directed reachability by proving that it is in fact not in monadic NP. The current paper greatly simplifies their proof.

The nonexpressibility results in this paper also hold in the presence of a larger class of built-in relations than those allowed in [1, 12]. There are three main reasons why built-in relations are important. First, proving nonexpressibility results for monadic NP in the presence of built-in relations seems to provide a good training ground for attacking the general (not monadic) case. For instance, before proving nonexpressibility results for "binary" NP, one clearly needs to deal with the case of monadic NP with a built-in successor or linear order relation, since the existence of such relations can be expressed by a binary second-order existential quantifier.

Another reason for interest in built-in relations is that there are known examples of classes for which built-in relations provably add to their expressive power. For example, the property "evenness" (i.e., the graph having an even number of nodes) is not in monadic NP, but it is in monadic NP with a built-in successor relation. Furthermore, some connections between computational complexity and descriptive complexity are known to hold only in the presence of built-in relations. As mentioned earlier, Immerman and Vardi showed that a property is in P iff it can be expressed in fixpoint logic with a built-in successor relation (or a built-in linear order). Allowing successor is crucial in this case, since evenness is not definable in fixpoint logic without successor [4].

Finally, built-in relations can be viewed as adding an element of nonuniformity to the class, and thus changing it somewhat (this is analogous to the way circuit-based computational complexity differs from Turing-machine based complexity). Proving that a problem is not in monadic NP even in the presence of certain built-in relations shows that the problem cannot even be captured in certain *nonuniform* ways (since the built-in relations vary from universe to universe), which is a stronger result.

The rest of the paper is organized as follows. In Section 2, we give definitions and conventions. In Section 3, we discuss first-order Ehrenfeucht–Fraïssé games. In Section 4, we give our new sufficient condition for the duplicator to have a winning strategy. In Section 5, we define monadic NP and Ajtai–Fagin games. In Section 6, we give our first application of our sufficient condition for the duplicator to have a winning strategy, by giving a simple proof of Fagin's result that connectivity is not in monadic NP (our proof is modeled after that of Fagin et al.). In Section 7, we consider the directed reachability problem. In Section 7.1, we sketch Ajtai and Fagin's proof [1] that directed reachability is not in monadic NP. In Section 7.2, we give our simplified proof of this result. This is our main application of our new sufficient condition for the duplicator to have a winning strategy. In Section 8, we give new inexpressibility results in the presence of certain built-in relations. We summarize in Section 9.

Other related work. A related development (independent of this paper) is a recent result by Schwentick [19]. He gives another sufficient condition for the duplicator to have a winning strategy, and uses it to show that connectivity is not in monadic NP in the presence of an even larger class of built-in relations than Fagin et al. had shown. Most importantly, he resolved an open problem of de Rougemont [6], by showing that connectivity is not in monadic NP even in the presence of a built-in linear order.

2. Definitions and conventions

A *language* $\mathscr{L}$ (sometimes called a *similarity type*, a *signature*, or a *vocabulary*) is a finite set $\{P_1,\ldots,P_p\}$ of relation symbols, each of which has an arity, along with a finite set $\{c_1,\ldots,c_q\}$ of constant symbols. An *$\mathscr{L}$-structure* (or *structure over $\mathscr{L}$*, or simply *structure*) is a set A (called the *universe*), along with a mapping associating a relation R_i over A with each $P_i \in \mathscr{L}$, where R_i has the same arity as P_i, for $1 \leqslant i \leqslant p$, and associating a member of A with each constant symbol $c_i \in \mathscr{L}$, for $1 \leqslant i \leqslant q$. We may call R_i the *interpretation of* P_i (and similarly for the constant symbols). If the point a is the interpretation of the constant symbol c_i, then we may say that a is *labeled* c_i. The structure is called *finite* if A is. Unless otherwise stated, throughout the rest of this paper we make the assumption that all structures we consider are finite. All our nonexpressibility results hold for infinite structures too, but the proof in the infinite case is known to be trivial using the Compactness Theorem. We use the usual Tarskian truth semantics to define what it means for a structure to *obey* or *satisfy* a sentence.

Our main application in this paper is to directed reachability. Therefore, we shall take some liberties with standard terminology, by (usually) taking a "graph" to mean a directed graph with two distinguished points, labeled s and t respectively. Thus, in this paper, a "graph" is a structure where the language consists of a single binary relation symbol and two constant symbols, s and t. We are also interested in "colored graphs", which are structures where the language consists of a single binary relation symbol and two constant symbols, as is the case for graphs, along with some finite number of unary relation symbols. If G is a colored graph, where the interpretations of the unary relation symbols in the language are $U_1, \ldots, U_k$, then by the *color* of a point a in the universe of G, we mean a description of which U_i's the point a is a member of. Thus, there are 2^k colors.

If $t = \langle x_1, \ldots, x_k \rangle$ is a tuple, define $[t]$ to be the set $\{x_1, \ldots, x_k\}$ of points that appear in t. We define the *hypergraph associated with structure* $\mathbf{A}$ to be a hypergraph (V, F) whose universe V is the same as the universe of $\mathbf{A}$ and whose set F of (hyper)edges is

$$\{[t] : t \text{ is a tuple in some relation of } \mathbf{A}\}.$$

A *(simple) path of length* k between two points u, v of $\mathbf{A}$ consists of a set of edges $S_1, \ldots, S_k \in F$ and a set of points $x_1, \ldots, x_{k-1} \in V$ such that (i) the x_i's are distinct from each other and from u and v, (ii) $u \in S_1$, (iii) $v \in S_k$, and (iv) $x_i \in S_i \cap S_{i+1}$, for $1 \leqslant i < k$. The *distance* between distinct points u and v is the smallest k such that there is a path of length k between them, and the distance between a point and itself is 0. If $k \geqslant 3$, then a *cycle of length* k in a structure $\mathbf{A}$ is a path of length k from a vertex to itself. (Shortly, we shall mention why we do not consider cycles of length 1 or 2.) Except for the fact that we do not consider cycles of length 2, this definition corresponds to Berge's notion [3] of a cycle in a hypergraph. (There are various other notions of a cycle in a hypergraph that are not equivalent to Berge's; see [11].) We note that if $\mathbf{A}$ is a structure over a language with a single binary relation, then its hypergraph is an ordinary undirected graph, and the concept of distance and cycle are the familiar ones.

We let $\langle x_1, x_2 \rangle$ represent the directed edge from x_1 to x_2 in a directed graph, and (x_1, x_2) the undirected edge between x_1 and x_2 in an undirected graph. By the *undirected version* of a directed graph, we mean the undirected graph obtained by ignoring the directions on the edges; thus, (x_1, x_2) is an edge in the undirected version iff either $\langle x_1, x_2 \rangle$ or $\langle x_2, x_1 \rangle$ is an edge in the directed graph. The *Gaifman graph* [14] of a structure $\mathbf{A}$ is the undirected graph with the same universe as $\mathbf{A}$, and with an edge (x_1, x_2) whenever x_1 and x_2 are distinct and appear together in a tuple of some relation of $\mathbf{A}$. In particular, the Gaifman graph of a directed graph without self-loops is simply the undirected version of the graph. Our definition of the distance between two points in a structure is equivalent to the distance between the two points in the Gaifman graph of the structure. However, the notions of a cycle in structure $\mathbf{A}$ and a cycle in the Gaifman graph of $\mathbf{A}$ are different in general. For example, if there is a tuple $\langle x_1, x_2, x_3 \rangle$ in a ternary relation of a structure $\mathbf{A}$ with all entries distinct, then there is a cycle of length 3 in the Gaifman graph (with edges (x_1, x_2), (x_2, x_3), and (x_3, x_1)), but not necessarily a cycle in $\mathbf{A}$. In general, a cycle in a structure gives rise to a cycle in the Gaifman graph, but not vice versa. Note that an assumption of our main theorem is that there are *no* small cycles in the structure. Thus, the fact that our notion of "cycle" is restrictive only increases the applicability of our theorem. This is also why we do not consider cycles of length less than 3: such very small cycles would have no effect on our theorems, and so we do not want to forbid them.

Finally, we define the *degree* of a point in $\mathbf{A}$ to be the degree in the Gaifman graph.

3. Ehrenfeucht–Fraïssé games

Most tools of model theory do not "survive" when we restrict our attention to finite structures. Ehrenfeucht–Fraïssé games [7, 13] are among the few that do. For an introduction to Ehrenfeucht–Fraïssé games and some of their applications to finite-model theory, see [1, pp. 122–126].

We begin with an informal definition of an *r-round first-order Ehrenfeucht–Fraïssé game* (where r is a positive integer), which we shall call an *r-game* for short. It is straightforward to give a formal definition, but we shall not do so. There are two *players*, called *the spoiler* and *the duplicator*, and two structures, $\mathbf{A}_0$ and $\mathbf{A}_1$. In the first round, the spoiler selects a point in one of the two structures, and the duplicator selects a point in the other structure. Let p_1 be the point selected in $\mathbf{A}_0$, and let q_1 be the point selected in $\mathbf{A}_1$. Then the second round begins, and again, the spoiler selects a point in one of the two structures, and the duplicator selects a point in the other structure. Let p_2 be the point selected in $\mathbf{A}_0$, and let q_2 be the point selected in $\mathbf{A}_1$. This continues for r rounds. The duplicator wins if the substructure of $\mathbf{A}_0$ induced by $p_1,\ldots,p_r$ is isomorphic to the substructure of $\mathbf{A}_1$ induced by $q_1,\ldots,q_r$, under the function that maps p_i onto q_i for $1 \leqslant i \leqslant r$. That is, the duplicator wins precisely if (a) $p_i = p_j$ iff $q_i = q_j$, for each i and j; (b) $\langle p_{i_1},\ldots,p_{i_\ell}\rangle$ is a tuple in a relation in $\mathbf{A}_0$ iff $\langle q_{i_1},\ldots,q_{i_\ell}\rangle$ is a tuple in the corresponding relation in $\mathbf{A}_1$, for each choice of $i_1,\ldots,i_\ell$; and (c) p_i is the interpretation in $\mathbf{A}_0$ of a constant symbol d iff q_i is the interpretation in $\mathbf{A}_1$ of the constant symbol d, for each i. Otherwise, the spoiler wins. We say that the spoiler or the duplicator *has a winning strategy* if he can guarantee that he will win, no matter how the other player plays. Since the game is finite, and there are no ties, the spoiler has a winning strategy iff the duplicator does not. If the duplicator has a winning strategy, then we write $\mathbf{A}_0 \sim_r \mathbf{A}_1$. In this case, intuitively, $\mathbf{A}_0$ and $\mathbf{A}_1$ are indistinguishable by an r-game.

The following important theorem (from [7, 13]) shows why these games are of interest. If $\mathscr{S}$ is a class of structures over a language $\mathscr{L}$, then let $\overline{\mathscr{S}}$ be the complement of $\mathscr{S}$, that is, the class of structures over $\mathscr{L}$ not in $\mathscr{S}$.

Theorem 3.1. *$\mathscr{S}$ is first-order definable iff there is r such that whenever $\mathbf{A}_0 \in \mathscr{S}$ and $\mathbf{A}_1 \in \overline{\mathscr{S}}$, then the spoiler has a winning strategy in the r-game over $\mathbf{A}_0, \mathbf{A}_1$.*

4. A winning condition for the duplicator

According to Theorem 3.1, to show that a property of finite structures is not first-order definable, we have to construct for each r a structure $\mathbf{A}_0$ satisfying the property and a structure $\mathbf{A}_1$ failing the property, such that the duplicator has a winning strategy in the r-game over $\mathbf{A}_0, \mathbf{A}_1$. Finding such structures can be nontrivial. The pair $\mathbf{A}_0, \mathbf{A}_1$ are guaranteed to look different at a "global" level, since $\mathbf{A}_0$ satisfies the property and $\mathbf{A}_1$ does not. How can they look similar in the Ehrenfeucht–Fraïssé game? The main observation (which has been made many times before) is that an Ehrenfeucht–Fraïssé game in some sense involves only "local" properties of the structures.

Next, we describe a condition which, if satisfied by $\mathbf{A}_0, \mathbf{A}_1$, guarantees that the duplicator has a winning strategy in an r-game. At an intuitive level, the condition says that the two structures are approximately isomorphic at a "local" level. As we will see, this approximate isomorphism leaves plenty of room for the structures to differ at the global level, and therein lies the condition's usefulness: a researcher using it to prove nonexpressibility results gets that much more "room" to construct the desired $\mathbf{A}_0, \mathbf{A}_1$. For example, in our proof that directed reachability is not in monadic NP, a probabilistic method suffices to prove the existence of the two structures. (This use of the probabilistic method is derived from the techniques of [1].)

Let d and q be integers, and let $\mathbf{A}$ be a structure. We define the notion of the *(d,q)-color* of each vertex in $\mathbf{A}$. Intuitively speaking, the (d,q)-color describes a small neighborhood around the vertex. For simplicity, we begin first by considering the case where $\mathbf{A}$ is a colored graph (this is the case of interest in the proof that directed reachability is not in monadic NP). Define the *$(d,0)$-color* of a vertex v to be the color in the colored graph, along with a description of whether or not there is an edge (a "self-loop") from the vertex v to itself, whether or not the vertex v is the distinguished point labeled s, and whether or not the vertex v is the distinguished point labeled t. Inductively, define the *$(d,q+1)$-color* of the vertex v (where $q \geqslant 0$) to be (a) a description of its (d,q)-color, along with (b) a complete description, for each possible (d,q)-color τ, as to whether there are $0,1,\ldots,d-1$, or at least d points w with (d,q)-color τ such that $\langle v,w\rangle$ is an edge of the graph, but $\langle w,v\rangle$ is not an edge, (c) a complete description, for each possible (d,q)-color τ, as to whether there are $0,1,\ldots,d-1$, or at least d points w with (d,q)-color τ such that $\langle w,v\rangle$ is an edge but $\langle v,w\rangle$ is not an edge, and (d) a complete description, for each possible (d,q)-color τ, as to whether there are $0,1,\ldots,d-1$, or at least d points w with (d,q)-color τ such that $\langle v,w\rangle$ and $\langle w,v\rangle$ are each edges. Thus, the $(d,q+1)$-color of a vertex v tells the (d,q)-color of v, and also tells how many vertices there are of each (d,q)-color with just an outedge from v, just an inedge into v, and both an outedge from and an inedge into v, where in all cases we do not count beyond d.

Readers familiar with [1] will recall that a slightly different notion of (d,q)-color appears there. That notion involves additional information, since under that definition, the $(d,q+1)$-color of a vertex v also contains a complete description, for each possible (d,q)-color τ, as to whether there are $0,1,\ldots,d-1$, or at least d points w with (d,q)-color τ such that *neither* $\langle v,w\rangle$ nor $\langle w,v\rangle$ is an edge. Our notion seems more useful because it is completely local: only "nearby" points (points whose distance from v is at most q) affect the (d,q)-color of v.

We now discuss how to define the (d,q)-color of each vertex v in a structure $\mathbf{A}$ over an arbitrary language $\mathscr{L}$. We begin with a preliminary notion. An *m-type* (among the m variables $x_1,\ldots,x_m$) is a conjunction such that (a) for each i and j between 1 and m, exactly one of $x_i = x_j$ or $x_i \neq x_j$ is a conjunct, and (b) for each arity ℓ, each relation symbol $P \in \mathscr{L}$ of arity ℓ, and each choice of $i_1,\ldots,i_\ell$ where $1 \leqslant i_j \leqslant m$ for each j, exactly one of $Px_{i_1}\dots x_{i_\ell}$ or $\neg Px_{x_1}\dots x_{i_\ell}$ is a conjunct. Intuitively, an m-type tells exactly how the variables $x_1,\dots,x_m$ relate to each other in a quantifier-free way. We say that the variable x *has a positive occurrence* in the m-type F if $Px_{i_1}\dots x_{i_\ell}$ (as opposed to $\neg Px_{i_1}\dots x_{i_\ell}$) is a conjunct of F for some relation symbol $P\in\mathscr{L}$ and some variables $x_{i_1},\dots,x_{i_\ell}$ where $x\in\{x_{i_1},\dots,x_{i_\ell}\}$. We define an m-type of vertices (as opposed to variables) analogously. Specifically, if $v_1,\dots,v_m$ are m vertices of $\mathbf{A}$, then we define their m-type to be an m-type among m variables $x_1,\dots,x_m$ that holds in $\mathbf{A}$ when $x_1,\dots,x_m$ are interpreted by $v_1,\dots,v_m$ respectively. Intuitively, an m-type among the m vertices $v_1,\dots,v_m$ of $\mathbf{A}$ tells exactly how these vertices relate to each other in $\mathbf{A}$. Similarly, we define what it means for a vertex v to have a positive occurrence in an m-type.

We are now ready to define the (d,q)-color of each vertex v in a structure $\mathbf{A}$ over an arbitrary language $\mathscr{L}$. Let m be the largest arity among relation symbols of $\mathscr{L}$. The $(d,0)$-color of v is a complete description of which relations of $\mathbf{A}$ have the tuple $\langle v,\dots,v\rangle$ as a member (where, of course, the length of the tuple is the arity of the relation), and which constant symbols label v. Inductively, define the *$(d,q+1)$-color* of the vertex v (where $q\geqslant 0$) to be a description of its (d,q)-color, along with a complete description, for each possible choice $\tau_1,\dots,\tau_{m-1}$ of (d,q)-colors and each possible m-type F among m vertices where v has a positive occurrence, as to whether there are $0,1,\dots,d-1$, or at least d choices of $\langle v_1,\dots,v_{m-1}\rangle$ such that $v, v_1,\dots,v_{m-1}$ have the m-type F and v_i has (d,q)-color τ_i for $1\leqslant i\leqslant m-1$. Finally, the (d,q) color of a tuple $\langle x_1,\dots,x_k\rangle$ is the tuple $\langle\tau_1,\dots,\tau_k\rangle$, where τ_i is the (d,q)-color of x_i.

The *multiplicity* of the (d,q)-color of a vertex is the number of vertices in the structure with this (d,q)-color. Similarly, the multiplicity of the (d,q)-color of a tuple in a relation in a structure is the number of tuples in that relation with this (d,q)-color.

The next theorem describes the desired sufficient condition.

Theorem 4.1. *Let r,f be positive integers. There is a positive integer k that depends only on r such that $\mathbf{A}_0\sim_r\mathbf{A}_1$ whenever $\mathbf{A}_0$ and $\mathbf{A}_1$ are structures of the same similarity type that have the same set of vertices and that satisfy the following conditions:*
1. *the degree of every vertex in $\mathbf{A}_0$ or $\mathbf{A}_1$ is at most f;*
2. *there is no cycle in either $\mathbf{A}_0$ or $\mathbf{A}_1$ of length less than k;*
3. *each vertex has the same (r,k)-color in $\mathbf{A}_0$ as in $\mathbf{A}_1$; and*
4. *if e is a tuple that is present in some relation in one structure but not in the corresponding relation in the other structure, then there are at least f^k tuples in both of these relations that have the same (r,k)-color as e.*

Remark. In Section 8 we will give a version of this theorem in which small cycles are allowed, but there are some additional assumptions.

We shall prove this theorem shortly. This theorem gives a sufficient condition for the duplicator to have a winning strategy. As we noted, Fagin et al. [12] make use of another sufficient condition, due essentially to Hanf, for the duplicator to have a winning strategy. Hanf's condition is incomparable with our condition. In particular, to apply Hanf's condition, it is necessary that for each point u in either structure there is a point v in the other structure such that u and v have isomorphic neighborhoods. If, however, $\mathbf{A}_0$ and $\mathbf{A}_1$ are directed graphs that differ only in that there is an edge e of $\mathbf{A}_0$ that is not in $\mathbf{A}_1$ (which, as we shall see, is how we use Theorem 4.1 to prove that directed connectivity is not in monadic NP), then Hanf's condition cannot be applied. This is because a neighborhood in $\mathbf{A}_0$ of an endpoint of e might not be isomorphic to a neighborhood of any point in $\mathbf{A}_1$. Intuitively, instead of requiring isomorphic neighborhoods (which Hanf's condition demands), our condition requires only "approximately isomorphic neighborhoods" (by dealing only with (r,k)-colors, rather than isomorphism types), but at the expense of adding other requirements (such as that there be no small cycles).

Although our Theorem 4.1 does not subsume Hanf's condition, nevertheless Theorem 4.1 is strong enough to replace Hanf's condition in Fagin et al.'s proof that connectivity is not in monadic NP (including the case of the built-in relations that they consider). This helps show the "versatility" of Theorem 4.1. In Section 6, we shall give our proof, using Theorem 4.1, that connectivity is not in monadic NP.

We conclude this section by giving a proof of Theorem 4.1. For ease in description, we prove this theorem under the assumption that $\mathbf{A}_0$ and $\mathbf{A}_1$ are colored graphs, rather than structures of arbitrary similarity type. (As we noted before, this is the case of interest in the proof that directed reachability is not in monadic NP.) It is fairly straightforward to modify our proof to deal with the general case, by replacing "distance in graph" with "distance in hypergraph," and other such changes.

Assuming the hypothesis of the theorem, we will describe how the duplicator matches the spoiler's moves, such that the substructures of $\mathbf{A}_0, \mathbf{A}_1$ picked at the end of r rounds are isomorphic. Before the formal proof, we give some intuition as to how we make use of (r,k)-colors, and why we do not allow small cycles.

Assume that $j < r \leqslant k$, that the spoiler and duplicator are playing an r-game on the colored graphs $\mathbf{A}_0$ and $\mathbf{A}_1$, and that through the first j rounds, the points $p_1, \dots, p_j$ picked in $\mathbf{A}_0$ have the same $(r-j,k)$-colors as the corresponding points $q_1, \dots, q_j$ picked in $\mathbf{A}_1$. Assume that the spoiler then picks, say, p_{j+1} in $\mathbf{A}_0$ different from $p_1, \dots, p_j$, where p_{j+1} is adjacent to p_i in $\mathbf{A}_0$ for some $i \leqslant j$ (perhaps, say, $\langle p_i, p_{j+1}\rangle$ is an edge but $\langle p_{j+1}, p_i\rangle$ is not). Since p_i and q_i have the same $(r-j,k)$-color, it follows that there is a point q_{j+1} in $\mathbf{A}_1$ different from $q_1, \dots, q_j$ that is adjacent to q_i in $\mathbf{A}_1$ (where, as before, $\langle q_i, q_{j+1}\rangle$ is an edge but $\langle q_{j+1}, q_i\rangle$ is not). If we assume that there are no cycles of length less than k in either $\mathbf{A}_0$ or $\mathbf{A}_1$, then p_{j+1} is not adjacent to any $p_m \neq p_i$, and q_{j+1} is not adjacent to any $q_m \neq q_i$. This enables us to maintain an isomorphism between the subgraph of $\mathbf{A}_0$ induced by the p_i's and the subgraph of $\mathbf{A}_1$ induced by the q_i's. If small cycles were possible, then it could happen that q_{j+1} would be adjacent to $q_m \neq q_i$ without p_{j+1} being adjacent to p_m, and this would violate isomorphism, and so the spoiler would win. This is why we do not allow small cycles. The full winning strategy for the duplicator is more complicated than this sketch (for example, we did not discuss what happens when the spoiler selects a point not adjacent to any point that has been selected so far); we now fill in the details.

Define $d(u,v)$ to be the distance between u and v, where as before, we ignore the directions of the edges. For ease in notation, we are a little sloppy by using $d(\cdot,\cdot)$ for both $\mathbf{A}_0$ and $\mathbf{A}_1$, even though distances may be different in them. It should always be clear in context which graph the distance is being measured in. We use the shorthand $u \rightsquigarrow v$ to denote the induced directed subgraph whose undirected version is a shortest path between u and v (this path will be unique wherever we use this notation). We will use the word "path" to refer either to an undirected path, or to the directed graph $u \rightsquigarrow v$, whose undirected version is a path. Define $d_m(u,v)$ to be $\min\{m, d(u,v)\}$ (i.e. "if the distance is at least m, then we do not care about the exact distance, but only that it is at least m"). Let $\mathrm{Ball}(v,q)$ be the set of vertices of distance at most q from v. We say that a directed graph (possibly colored) is a *tree* if it is a tree when we ignore the directions on the edges. Thus, it must be connected and have no (undirected) cycles.

As a part of our proof of Theorem 4.1, we state and prove three facts and a lemma. The following definition will be useful. Let X be a set of vertices in graph G. Define $W(X,K)$ to consist of X, along with all vertices that appear in a path of length at most K between some pair of points in X. Let $\mathrm{sg}_G(X,K)$ denote the subgraph of G induced by the set $W(X,K)$ of vertices. Note that even though we ignore edge directions when we consider the paths, $\mathrm{sg}_G(X,K)$ itself is a directed graph (the directions of the edges are determined by G). If $\mathbf{A}$ is a colored graph, then $\mathrm{sg}_{\mathbf{A}}(X,K)$ is the colored subgraph of $\mathbf{A}$ that is defined similarly.

The first fact says that in graphs with no small cycles, small neighborhoods around a point look like trees.

Fact 1. *Assume that a graph G has no cycles of length less than k (where $k \geqslant 4$), and $\{u_1, u_2, \ldots, u_m\}$ is a set of vertices. Assume that there is some ℓ with $1 \leqslant \ell \leqslant m$ such that $d(u_i, u_\ell) \leqslant k/6$ for each i with $1 \leqslant i \leqslant m$. Then $\mathrm{sg}_G(\{u_1, \ldots, u_m\}, k/3)$ is a tree.*

Proof. Let Y be the set consisting of $\{u_1, u_2, \ldots, u_m\}$, along with, for every i with $1 \leqslant i \leqslant m$, all vertices on a shortest (undirected) path between u_ℓ and u_i. Let T be the subgraph of G induced by the set Y of vertices. We now show that T is a tree. It is clearly connected. Assume that T contains a cycle; let C be a cycle of minimum size in T. Since T is an induced subgraph of G, we know by assumption that the cycle C has length at least k. Therefore, there are vertices x, y on the cycle with a path P_1 between them of length $\lfloor k/2 \rfloor$. Since $d(u_\ell, x) \leqslant k/6$ and $d(u_\ell, y) \leqslant k/6$, there is a path P_2 between x and y of length at most $k/3$. Now $k/3 < \lfloor k/2 \rfloor$, since $k \geqslant 4$. Therefore, the paths P_1 and P_2 are distinct. So there are two distinct paths between x and y, one path of length at most $k/3$, and the other of length $\lfloor k/2 \rfloor$. It follows that there is a cycle of length at most $(k/3) + \lfloor k/2 \rfloor < k$. This is a contradiction. Therefore, T is indeed a tree.

We now show that $sg_G(\{u_1,\dots,u_m\},k/3)$ is a tree. It is clearly connected, since $d(u_\ell,u_i)\leqslant k/6$ for each i. Since T is a tree, we need only show that for each i,j, each path of length at most $k/3$ between u_i and u_j is a path in T. If not, then there is a path in G of length at most $k/3$ between u_i and u_j that is not a path in T. Since $d(u_\ell,u_i)\leqslant k/6$ and $d(u_\ell,u_j)\leqslant k/6$, there is a path in T between u_i and u_j of length at most $k/3$. So there are two distinct paths in G between u_i and u_j (one path not in T, and one path in T), each of length at most $k/3$. It follows that there is a cycle of length at most $2k/3 < k$ in G. This is a contradiction. □

The next observation will be used in showing that the duplicator can always match the spoiler's actions. To motivate it, we remind the reader that the (d,q)-color of a vertex u is a partial description of the colored Ball(u,q). For instance, if u,v are vertices of the same (d,q)-color and there is vertex at distance x from u whose $(d,q-x)$-color is τ, then there is a vertex of the same $(d,q-x)$-color at distance x from v. Now we generalize this.

Fact 2. *Assume that colored graphs $\mathbf{A}_0$ and $\mathbf{A}_1$ have no cycles of length less than k, and m,r are integers such that $m<r$. Let $\{p_1,p_2,\dots,p_m\}$ and $\{q_1,\dots,q_m\}$ be subsets of vertices of $\mathbf{A}_0$ and $\mathbf{A}_1$, respectively, such that for some ℓ, where $1\leqslant\ell\leqslant m$, $\{p_1,p_2,\dots,p_m\}\subseteq Ball(p_\ell,k/6)$ and $\{q_1,\dots,q_m\}\subseteq Ball(q_\ell,k/6)$. Assume also that $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m\},k/3)$ and $sg_{\mathbf{A}_1}(\{q_1,\dots,q_m\},k/3)$ are isomorphic as colored trees, under the isomorphism mapping p_i to q_i for each i, where each vertex is colored with its (r,y)-color. Let v be a vertex in $\mathbf{A}_0$, and let $x=d(v,p_\ell)$. Assume that $x\leqslant k/6$ and $x\leqslant y$. Then there is a point u in $\mathbf{A}_1$ such that $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m,v\},k/3)$ and $sg_{\mathbf{A}_1}(\{q_1,\dots,q_m,u\},k/3)$ are isomorphic as colored trees, under the isomorphism mapping p_i to q_i for each i and mapping v to u, where each vertex is colored with its $(r,y-x)$-color.*

Proof. The fact that all the graphs of the form $sg_{\mathbf{A}_j}(X,k/3)$ mentioned above are necessarily trees follows from Fact 1. If $v\in sg_{\mathbf{A}_0}(\{p_1,\dots,p_m\},k/3)$, then the corresponding vertex in $sg_{\mathbf{A}_1}(\{q_1,\dots,q_m\},k/3)$ satisfies all the conditions required for u (and in fact has the same (r,y)-color as v, and not just the same $(r,y-x)$-color). So assume $v\notin sg_{\mathbf{A}_0}(\{p_1,\dots,p_m\},k/3)$, and let i be such that p_i is the vertex in $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m\},k/3)$ closest to v. Since $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m,v\},k/3)$ is a tree, it follows that for all $t\leqslant m$, the path $v\rightsquigarrow p_t$ must pass through p_i. So $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m,v\},k/3)$ is the union of $sg_{\mathbf{A}_0}(\{p_1,\dots,p_m\},k/3)$ and the path $p_i\rightsquigarrow v$. We have to show that an isomorphic path can be added to $sg_{\mathbf{A}_1}(\{q_1,\dots,q_m\},k/3)$.

Let $v_1,\dots,v_{j+1}$ be the sequence of vertices on the path $p_i\rightsquigarrow v$, where $v_1=p_i$, $v_{j+1}=v$ and $j\leqslant x$. Let τ be the $(r,y-1)$-color of v_2. Let V be the set of vertices v in $\mathbf{A}_0$ with $(r,y-1)$-color τ that relate to v_1 in the same way as v_2 relates to v_1; that is, $\langle v_1,v\rangle$ (resp., $\langle v,v_1\rangle$) is a directed edge of $\mathbf{A}_0$ iff $\langle v_1,v_2\rangle$ (resp., $\langle v_2,v_1\rangle$) is a directed edge of $\mathbf{A}_0$. Similarly, let U be the set of vertices in $\mathbf{A}_1$ with $(r,y-1)$-color τ that relate to u_1 in the same way as v_2 relates to v_1. Since q_i and p_i have the same (r,y)

color, it follows that either U and V have the same cardinality, or both cardinalities are at least r. Since V contains a vertex not in $\{p_1,\ldots,p_m\}$, and since $m < r$, it follows that U contains a vertex not in $\{q_1,\ldots,q_m\}$. Let u_2 be such a vertex. We can now argue similarly about the neighbors of v_2 and u_2 (but where we do not need to depend on the fact that $m < r$) to find a vertex u_3, and keep doing this until we have found a path outside of $\{q_1,\ldots,q_m\}$ isomorphic to $p_i \rightsquigarrow v$, when vertices are colored with their $(r, y-j)$-color. Define u to be the final vertex on the path. $\square$

Finally, since $\mathrm{Ball}(v,a)$ has at most f^{a+1} vertices in the graphs of Theorem 4.1, we have the following facts about vertex colors of high multiplicity.

Fact 3. *Let f be the maximum degree of the vertices in a graph.*
1. *Let τ be a vertex color of multiplicity more than $i \cdot f^{a+1}$. Given any i points $p_1,\ldots,p_i$, there is a vertex of color τ outside $\bigcup_{\ell \leqslant i} \mathrm{Ball}(p_\ell, a)$.*
2. *Let u,v be vertices such that $d(u,v) = x$. If the multiplicity of the (d,q)-color of v is m, then the multiplicity of the $(d,q-x)$-color of u is at least m/f^x.*

Proof. The first part is obvious. We now prove the second part. Note first that the number of vertices of distance exactly x from a given vertex is at most f^x. Denote by τ the (d,q)-color of v and by τ' the $(d,q-x)$-color of u. Then each of the m vertices having color τ must have a vertex of color τ' at a distance x from it. This yields m vertices of color τ', but where we have counted each vertex w of color τ' at most f^x times (once for each vertex of (d,q)-color τ of distance exactly x from w). So there are at least m/f^x vertices of $(d,q-x)$-color τ'. $\square$

Now we are ready to describe the duplicator's strategy and prove Theorem 4.1. If the maximal degree f is 1, then the result follows easily. So assume that $f \geqslant 2$. Let $k = 3^{2r}$. The strategy maintains the following invariant.

Invariant after round i. Sequences $p_1, p_2, \ldots, p_i$ and $q_1, q_2, \ldots, q_i$ of vertices have been picked in $\mathbf{A}_0$ and $\mathbf{A}_1$ respectively.
- *Isomorphism invariant*: $\mathrm{sg}_{\mathbf{A}_0}(\{p_1,\ldots,p_i\}, 3^{r-i})$ is isomorphic to $\mathrm{sg}_{\mathbf{A}_1}(\{q_1,\ldots,q_i\}, 3^{r-i})$ under the isomorphism mapping p_j to q_j for each $j \leqslant i$, when the vertices are colored with their $(r, 3^{2r-i})$-color.
- *Multiplicity invariant*: If $j \leqslant i$ is such that $p_j \neq q_j$, then the multiplicity of their common $(r, 3^{2r-i})$-color is at least $f^{3^{r+1}-3^r-3^{r-1}-\cdots-3^{r-i}}$.

Then the following generalization of Fact 2 shows that the isomorphism invariant is easy to maintain if the next vertex picked by the spoiler is in $\bigcup_{j \leqslant i} \mathrm{Ball}(p_j, 3^{r-i-1})$.

Lemma 4.2. *For each vertex $v \in \bigcup_{\ell \leqslant i} \mathrm{Ball}(p_\ell, 3^{r-i-1})$, there is a vertex $u \in \bigcup_{\ell \leqslant i} \mathrm{Ball}(q_\ell, 3^{r-i-1})$ such that the isomorphism invariant is maintained for $p_1,\ldots,p_i,v$ and $q_1,\ldots,q_i,u$, where $i+1$ plays the role of i.*

Proof. Let ℓ be such that $v \in \mathrm{Ball}(p_\ell, 3^{r-i-1})$. Let $\{p_{\ell_1}, \dots, p_{\ell_m}\}$ be the subset of p_i's lying in $\mathrm{Ball}(p_\ell, 3^{r-i})$. Then $q_{\ell_1}, \dots, q_{\ell_m}$ must lie in $\mathrm{Ball}(q_\ell, 3^{r-i})$. Furthermore, $\mathrm{sg}_{\mathbf{A}_0}(\{p_{\ell_1}, \dots, p_{\ell_m}\}, 3^{r-i})$ and $\mathrm{sg}_{\mathbf{A}_1}(\{q_{\ell_1}, \dots, q_{\ell_m}\}, 3^{r-i})$ must be isomorphic colored trees, under the isomorphism mapping p_{ℓ_i} to q_{ℓ_i} for each $i \leqslant m$, when the vertices are colored with their $(r, 3^{2r-i})$-color. Now Fact 2 can be applied to yield a vertex $u \in \mathrm{Ball}(q_\ell, 3^{r-i-1})$ with the same $(r, 3^{2r-i} - 3^{r-i-1})$-color as v, and such that the isomorphism invariant is maintained for the trees formed by $p_{\ell_1}, \dots, p_{\ell_m}, v$ and $q_{\ell_1}, \dots, q_{\ell_m}, u$ using the $(r, 3^{2r-i} - 3^{r-i-1})$-color. Since $3^{2r-i} - 3^{r-i-1} \geqslant 3^{2r-i-1}$ for $i < r$, we conclude that the trees are isomorphic for the $(r, 3^{2r-i-1})$-color too.

To finish the claim about the isomorphism invariant, it suffices to show that the distance of u (resp., v) to a vertex p_t outside $\mathrm{Ball}(q_\ell, 3^{r-i})$ (resp., vertex q_t outside $\mathrm{Ball}(p_\ell, 3^{r-i})$) is more than 3^{r-i-1}, so that those vertices do not interfere with the isomorphism invariant. But this follows from the triangle inequality, since $d(v, p_t) \geqslant d(p_\ell, p_t) - d(v, p_\ell) > 3^{r-i} - 3^{r-i-1} = 2 \cdot 3^{r-i-1}$. The same argument works for v and q_t. $\square$

We now describe the duplicator's strategy. Suppose the spoiler picks a vertex p_{i+1} in $\mathbf{A}_0$ in round $i+1$ (a symmetric strategy is used if the spoiler chooses a vertex in $\mathbf{A}_1$). Let u be the same vertex in $\mathbf{A}_1$. The conditions of Theorem 4.1 imply that u and p_{i+1} have the same $(r, 3^{2r})$-color. So they also have the same $(r, 3^{2r-i-1})$-color, say τ.

If letting q_{i+1} be u does not lead to a violation of the isomorphism invariant when $i+1$ plays the role of i, just do that. Otherwise there is a $j \leqslant i$ such that either $d_{3^{r-i-1}}(u, q_j) \neq d_{3^{r-i-1}}(p_{i+1}, p_j)$, or $d_{3^{r-i-1}}(u, q_j) = d_{3^{r-i-1}}(p_{i+1}, p_j) < 3^{r-i-1}$ and the colored path $u \rightsquigarrow q_j$ is not isomorphic to the colored path $p_{i+1} \rightsquigarrow p_j$ under the isomorphism mapping u to p_{i+1} and q_j to p_j, when the vertices are colored with their $(r, 3^{2r-i-1})$-color. In each case we show how to find a vertex maintaining the invariant.

Case 1: $d_{3^{r-i-1}}(u, q_j) \neq d_{3^{r-i-1}}(p_{i+1}, p_j)$ *and* $p_{i+1} \notin \bigcup_{\ell \leqslant i} \mathrm{Ball}(p_\ell, 3^{r-i-1})$. Since $p_{i+1} \notin \bigcup_{\ell \leqslant i} \mathrm{Ball}(p_\ell, 3^{r-i-1})$, it follows that $d_{3^{r-i-1}}(p_{i+1}, p_j) = 3^{r-i-1}$. Since also $d_{3^{r-i-1}}(u, q_j) \neq d_{3^{r-i-1}}(p_{i+1}, p_j)$, it follows that $u \in \mathrm{Ball}(q_j, 3^{r-i-1})$. We need a vertex of color τ not in $\bigcup_{\ell \leqslant i} \mathrm{Ball}(q_\ell, 3^{r-i-1})$. Such a vertex exists (by part 1 of Fact 3), if the multiplicity of τ is at least $i \cdot f^{3^{r-i-1}}$. We will show that the multiplicity of τ is at least $f^{3^{r+1} - 3^r - \dots - 3^{r-i-1}} \geqslant i \cdot f^{3^{r-i-1}}$, so the multiplicity invariant will also be maintained. We show this high multiplicity by analyzing the two possible cases: either $p_j = q_j$ or $p_j \neq q_j$. If $p_j = q_j$, then there is an edge on the path $u \rightsquigarrow q_j$ in $\mathbf{A}_1$ that is not in $\mathbf{A}_0$. By assumption, the multiplicity of the $(r, 3^{2r})$-color of the missing edge is at least $f^{3^{2r}}$. Since u is within a distance 3^{r-i-1} of this edge, part 2 of Fact 3 implies that the multiplicity of τ is at least $f^{3^{2r} - 3^{r-i-1}}$, which is at least $f^{3^{r+1} - 3^{r-i-1}}$, which in turn is at least $f^{3^{r+1} - 3^r - \dots - 3^{r-i-1}}$, as desired. Now assume that $p_j \neq q_j$. The existing multiplicity invariant implies that the multiplicity of the $(r, 3^{2r-i})$-color of p_j and q_j is at least $f^{3^{r+1} - 3^r - 3^{r-1} - \dots - 3^{r-i}}$. Since u is within distance 3^{r-i-1} of q_j, it follows from part 2 of Fact 3 that the multiplicity of the $(r, 3^{2r-i} - 3^{r-i-1})$-color of u is at least $f^{3^{r+1}-3^r-3^{r-1}-\cdots-3^{r-i}-3^{r-i-1}}$, as desired.

Case 2: $d_{3^{r-i-1}}(u, q_j) \neq d_{3^{r-i-1}}(p_{i+1}, p_j)$ *and* $p_{i+1} \in \bigcup_{\ell \leqslant i} \mathrm{Ball}(p_\ell, 3^{r-i-1})$. Then by using Lemma 4.2 we can find a point q_{i+1} in $\bigcup_{\ell \leqslant i} \mathrm{Ball}(q_\ell, 3^{r-i-1})$ while maintaining the isomorphism invariant. It remains to be shown that the multiplicity invariant is also maintained. Once again, there are two subcases: either $p_j = q_j$ or $p_j \neq q_j$. The analysis is similar to that in Case 1.

Case 3: $d_{3^{r-i-1}}(u, q_j) = d_{3^{r-i-1}}(p_{i+1}, p_j) < 3^{r-i-1}$ and the colored path $u \rightsquigarrow q_j$ is not isomorphic to the colored path $p_{i+1} \rightsquigarrow p_j$ under the isomorphism mapping u to p_{i+1} and q_j to p_j, when the vertices are colored with their $(r, 3^{2r-i-1})$-color. Then we can use Lemma 4.2 to find a q_{i+1} such that the isomorphism invariant is maintained. We only have to argue about the multiplicity invariant being maintained, since the procedure yields a q_{i+1} different from p_{i+1}. Again we have to consider the cases $p_j = q_j$ or $p_j \neq q_j$, and argue as in Case 1. This concludes the proof of Theorem 4.1.

5. Monadic NP and Ajtai–Fagin games

First-order logic allows for quantification only over members of the universe, and not over sets of members of the universe, or more generally, over relations; for details, see Enderton [8] or Shoenfield [20]. When we pass from first-order logic to second-order logic, we allow quantification over relations. In particular, a Σ^1_1 *sentence* is a sentence of the form $\exists A_1 \ldots \exists A_k \psi$, where ψ is first-order and where the A_i's are relation symbols. We now give three examples of Σ^1_1 sentences. In each of these examples, E represents the edge relation of the graph.

Example 5.1. We first construct a Σ^1_1 sentence that says that a graph has a Hamiltonian path (that is, the graph has a path that passes through all of the points). Let $<$ be a new binary relation symbol, and let ψ_1 say "$<$ is a linear order". Thus, ψ_1 is a conjunction of the following two sentences, the first of which says that $<$ is transitive, and the second of which says that $<$ satisfies trichotomy. In the second sentence, for convenience, we use $\oplus$ to represent "exclusive or".

$$\forall x \forall y \forall z((x < y) \wedge (y < z) \Rightarrow (x < z)).$$
$$\forall x \forall y((x < y) \oplus (x = y) \oplus (y < x)).$$

Let ψ_2 say "If y is the immediate successor of x, then there is an edge from x to y". Thus, ψ_2 is

$$\forall x \forall y(((x < y) \wedge \forall z \neg((x < z) \wedge (z < y))) \Rightarrow Exy).$$

The Σ^1_1 sentence $\exists < (\psi_1 \wedge \psi_2)$ then says "The graph has a Hamiltonian path".

Example 5.2. We now construct a Σ^1_1 sentence that says that a graph is 3-colorable. In this sentence, the three colors are represented by the unary relation symbols A_1, A_2, and A_3. Let ψ_1 say "Each point has exactly one color". Thus, ψ_1 is

$$\forall x((A_1x \wedge \neg A_2x \wedge \neg A_3x) \vee (\neg A_1x \wedge A_2x \wedge \neg A_3x) \vee (\neg A_1x \wedge \neg A_2x \wedge A_3x)).$$

Let ψ_2 say "No two vertices with the same color are connected by an edge". Thus, ψ_2 is

$$\forall x \forall y((A_1x \wedge A_1y \Rightarrow \neg Exy) \wedge (A_2x \wedge A_2y \Rightarrow \neg Exy) \wedge (A_3x \wedge A_3y \Rightarrow \neg Exy)).$$

The Σ_1^1 sentence $\exists A_1 \exists A_2 \exists A_3(\psi_1 \wedge \psi_2)$ then says "The graph is 3-colorable".

Example 5.3. Our final example, which is very relevant for this paper, deals with the class of graphs where there is no path from s to t. Let ψ_1 say "The set A contains s, and its complement contains t", that is, $As \wedge \neg At$. Let ψ_2 say "There is no edge from A to its complement", that is, $\forall x \forall y((Ax \wedge \neg Ay) \Rightarrow \neg Exy)$. It is clear that the Σ_1^1 sentence $\exists A(\psi_1 \wedge \psi_2)$ characterizes those graphs with no path from s to t.

A Σ_1^1 sentence $\exists A_1 \ldots \exists A_k \psi$, where ψ is first-order, is said to be *monadic* if each of the A_i's is unary, that is, the existential second-order quantifiers quantify only over sets. A class $\mathscr{S}$ of graphs is said to be (monadic) Σ_1^1 if it is the class of all finite graphs that obey some fixed (monadic) Σ_1^1 sentence. One reason that Σ_1^1 classes are of great interest is Fagin's result [9] that the collection of Σ_1^1 classes coincides with the complexity class NP. For this reason, we follow [12] and refer to the collection of monadic Σ_1^1 classes as *monadic NP*. The Σ_1^1 sentences constructed in Examples 5.2 and 5.3 are monadic, and so 3-colorability and nonreachability are in monadic NP. Note that 3-colorability is an NP-complete property [15]. Thus, monadic NP includes NP-complete properties. The Σ_1^1 sentence constructed in Example 5.1 is not monadic (since $<$ is a binary relation symbol). Indeed, Turán [21] has shown that Hamiltonicity is not in monadic NP; in fact, he showed the stronger result that Hamiltonicity cannot be defined by a monadic second-order sentence (where we allow arbitrary quantification, both universal and existential, over sets).

We now discuss a game that corresponds to monadic NP. This game, which is called the *Ajtai–Fagin (c,r)-game*, involves c colors and r rounds. It was introduced in [1] to prove that directed reachability is not in monadic NP. Let $\mathscr{S}$ be a class of graphs. For example, $\mathscr{S}$ could be the class of graphs that are (s,t)-connected. Let D be a set of c distinct colors. For simplicity, we can assume throughout this paper that $c = 2^k$ for some k, so that coloring a point x corresponds to deciding which of k unary relations x is a member of. The duplicator selects a member G_0 of $\mathscr{S}$. The spoiler then colors each of the points of G_0, using the colors in D, to obtain the colored graph $\mathbf{A}_0$. The duplicator then selects a member G_1 of $\overline{\mathscr{S}}$, the complement of $\mathscr{S}$. Then the duplicator colors each of the points of G_1, using the colors in D, to obtain the colored graph $\mathbf{A}_1$. Note that there is an asymmetry in the two graphs in the rules of the game, in that the spoiler must color the points of

G_0, not G_1. The game then concludes with an r-game. The duplicator now wins if he wins the final r-game. Ajtai and Fagin prove the following result, which is not hard.

Theorem 5.4. *$\mathscr{S}$ is in monadic* NP *iff there are c, r such that the spoiler has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$.*

6. Connectivity

In this section, we show how Theorem 4.1 can be used to prove that connectivity is not in monadic NP. This was first proved by Fagin [10], and then given a much simpler proof by Fagin et al. [12].

Theorem 6.1. *Connectivity is not in monadic* NP.

Proof. We follow the outline of Fagin et al.'s proof, but we use Theorem 4.1 instead of Hanf's condition.

Let c (the number of colors) and r (the number of rounds) be given. Let $\mathscr{S}$ be the class of connected graphs. By Theorem 5.4, we need only show that the duplicator has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$. As his first move in the Ajtai-Fagin (c,r)-game over $\mathscr{S}$, the duplicator selects the graph G_0 to be a directed cycle of length n, for a sufficiently large n (we shall discuss how large later). The spoiler then colors G_0 with the c colors, to obtain the colored graph $\mathbf{A}_0$. Let $\alpha_0, \alpha_1, \ldots, \alpha_{n-1}$ denote the points in order around the cycle, so that there is an edge $\langle \alpha_i, \alpha_{i+1} \rangle$ for $0 \leqslant i < n$ (we think of the subscripts as being reduced modulo n to belong to the interval $[0, n-1]$). Let $f = 2$ (the degree of each point on a cycle), and let k be as in Theorem 4.1. Since the number of (r,k)-colors is independent of the number n of points, it is easy to see that by picking n sufficiently large, there is a point α_p of $\mathbf{A}_0$ such that the multiplicity of the (r,k)-color τ of the edge $\langle \alpha_p, \alpha_{p+1} \rangle$ is at least f^k, and there is a point α_q of $\mathbf{A}_0$ that is of distance at least k from α_p such that the edge $\langle \alpha_q, \alpha_{q+1} \rangle$ has this same (r,k)-color τ. The duplicator then selects G_1 to be the graph obtained from G_0 by deleting the edges $\langle \alpha_p, \alpha_{p+1} \rangle$ and $\langle \alpha_q, \alpha_{q+1} \rangle$, and adding the edges $\langle \alpha_p, \alpha_{q+1} \rangle$ and $\langle \alpha_q, \alpha_{p+1} \rangle$. It is easy to see that G_1 is the disjoint union of two cycles, and in particular, is not connected. The duplicator then colors G_1, vertex by vertex, just as G_0 was colored, to obtain the colored graph $\mathbf{A}_1$. It is easy to see that each vertex v has the same (r,k)-color in $\mathbf{A}_0$ as in $\mathbf{A}_1$, since the colored Ball(v,k) in $\mathbf{A}_0$ and $\mathbf{A}_1$ are isomorphic. The edges that are present in $\mathbf{A}_0$ but not $\mathbf{A}_1$ are $\langle \alpha_p, \alpha_{p+1} \rangle$ and $\langle \alpha_q, \alpha_{q+1} \rangle$, which each have (r,k)-color τ. The edges that are present in $\mathbf{A}_1$ but not $\mathbf{A}_0$ are $\langle \alpha_p, \alpha_{q+1} \rangle$ and $\langle \alpha_q, \alpha_{p+1} \rangle$, which it is easy to see each have (r,k)-color τ. It then follows immediately from Theorem 4.1 and our choice of k that $\mathbf{A}_0 \sim_r \mathbf{A}_1$. Thus, the duplicator has a winning strategy in the remaining r-game. So the duplicator has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$, as desired. $\square$

7. Directed reachability

Our main application of Theorem 4.1 is to give a simpler proof of Ajtai and Fagin's result that directed reachability is not in monadic NP. We begin with a sketch of their proof.

7.1. A sketch of Ajtai and Fagin's proof

Ajtai and Fagin's proof makes use of the characterization of monadic NP from Theorem 5.4. Let $\mathscr{S}$ be the class of (s,t)-connected graphs. Let $v_1, \ldots, v_n$ be n points, which are used as the set of vertices of the graph G_0. The vertex v_1 is labeled s, and the vertex v_n is labeled t. Then G_0, the member of $\mathscr{S}$ selected by the duplicator in the (c,r)-game over $\mathscr{S}$, has "forward edges" $\langle v_i, v_{i+1} \rangle$ for $1 \leqslant i < n$; these form a path from s to t. In addition, G_0 has certain "backedges" $\langle v_i, v_j \rangle$ where $j < i$. The choice of these backedges are made by probabilistic means; it turns out that for the proof to work, there cannot be too few or too many backedges. We refer to such a graph G_0 as an *(s,t)-path with backedges.* If e is one of the forward edges of G_0, then denote by $G_0 - e$ the graph that results by deleting the edge e. It is clear that (a) there is a path from s to t in G_0, but (b) for each forward edge e, there is no path from s to t in $G_0 - e$. Thus, $G_0 \in \mathscr{S}$, but $G_0 - e \in \overline{\mathscr{S}}$ for each forward edge e. They now show that for a certain choice of G_0 and for each coloring of G_0 by the spoiler, there is a forward edge e such that the duplicator can select $G_1 \in \overline{\mathscr{S}}$ to be $G_0 - e$, color G_1 with exactly the same coloring, vertex by vertex, as G_0 is colored, and then have a winning strategy in the remaining r-game. By Theorem 5.4, this is sufficient to show that directed reachability is not in monadic NP.

The graph G_0 selected by the duplicator is guaranteed to exist by the next theorem, which the reader should be able to prove using standard arguments from the theory of random graphs.

Theorem 7.1. *Assume that c, d, q, and m are positive integers, and $\varepsilon > 0$. If n is sufficiently large, then there is a graph G_0 with n vertices that is an (s,t)-path with backedges such that*
1. *for every vertex v of G_0, the number of vertices whose distance in G_0 from v is at most m is less than n^ε;*
2. *the number of vertices that are on some cycle of size less than m in G_0 is less than n^ε; and*
3. *however the spoiler colors G_0 with the c colors, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that each vertex has the same (d,q)-color in G_0 as in $G_0 - e$.*

In part (3) of Theorem 7.1, the probability is given by taking all possible choices of the forward edge to be equally likely.

Of course, Ajtai and Fagin use their notion of (d,q)-color, which implies the result for our notion. Also, instead of n^ε in clauses (1) and (2), they have $n^{1/100}$ and $n^{3/4}$, respectively. This was just for convenience, and it is easy to see that their proof can give the result involving n^{ε}.

Theorem 7.1 tells us that given the parameters mentioned in the first sentence of the theorem, G_0 can be selected so that (1) every neighborhood of radius m is small, (2) there are few points on short cycles (cycles of size less than m), and (3) after the spoiler has colored G_0, then for almost all choices of the forward edge e, each vertex has the same (d,q)-color in G_0 as in $G_0 - e$. Theorem 7.1 is proven by selecting each backedge to appear with a certain probability, so that there are not too many and not too few backedges. The proof then proceeds by probabilistic arguments, which are not very difficult. The idea behind Theorem 7.1 is fairly intuitive. Since there are not too many backedges, it follows that neighborhoods are small, and the number of points on short cycles is small. This gives us (1) and (2) above. Certain forward edges are "special", in that, for example, they are near some point whose (d,q)-color is unusual, or they are near some point on a short cycle. Since neighborhoods are small, and the number of points on short cycles is small, it follows that nearly all forward edges are nonspecial. Then (3) above follows from the fact that there are enough backedges so that the absence of a nonspecial forward edge e is compensated for by the presence of many backedges. For example, if $e = \langle x, y \rangle$, then there are enough backedges from x to vertices with the same $(d,q-1)$-color as y that the edge e can be deleted without affecting the (d,q)-color of x. Here we take advantage of the fact that the (d,q)-color "counts only as high as d".

We now describe how Ajtai and Fagin use Theorem 7.1 to prove that directed reachability is not in monadic NP. Let c (the number of colors) and r (the number of rounds) be given. They want to show that the duplicator has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathcal{S}$, the class of directed graphs that are (s,t)-connected. They select d, q, m, and n to be sufficiently large with respect to c and r, and take $\varepsilon > 0$ sufficiently small. As his first move in the Ajtai–Fagin (c,r)-game over $\mathcal{S}$, the duplicator selects the graph G_0 guaranteed by Theorem 7.1, and the spoiler then colors G_0 with the c colors. Denote the resulting colored graph by $\mathbf{A}_0$. Let e be one of the forward edges guaranteed by part (3) of Theorem 7.1 so that each vertex has the same (d,q)-color in G_0 as in $G_0 - e$. The duplicator now selects $G_0 - e$ as a member of $\overline{\mathcal{S}}$, and colors $G_0 - e$, vertex by vertex, with the same coloring as G_0. Denote the resulting colored graph by $\mathbf{A}_1$. Ajtai and Fagin then give a complicated proof using elaborate combinatorial constructs to show that under a suitable choice of the parameters d, q, m, and ε, the duplicator has a winning strategy in the r-game played on $\mathbf{A}_0$ and $\mathbf{A}_1$. It then follows from Theorem 5.4 that directed reachability is not in monadic NP.

7.2. *Our simplified proof*

Our proof that directed reachability is not in monadic NP proceeds as follows. We make use of a slightly modified version of Theorem 7.1 (namely, Theorem 7.2 below), which is more useful for our purposes. The beginning of our proof is then similar to that of Ajtai and Fagin. However, we do not conclude the proof with their very complicated proof that the duplicator has a winning strategy on the r-game played over $\mathbf{A}_0$ and $\mathbf{A}_1$. Instead, we conclude the proof by using Theorem 4.1. The net result is an understandable and much simpler proof that directed reachability is not in monadic NP.

We now give the slight modification of Theorem 7.1 that we use in our proof. The graph G_0 guaranteed from Theorem 7.1 may have short cycles (cycles of size less than m), but there are few points on short cycles. It is convenient for us to completely eliminate the short cycles. Given G_0 as in Theorem 7.1, define G_0' by deleting some backedge on each short cycle. We thereby obtain the following theorem (where G_0 in the statement of the theorem below is G_0').

Theorem 7.2. *Assume that c, d, q, and m are positive integers , and $\varepsilon > 0$. If n is sufficiently large, then there is a graph G_0 with n vertices that is an (s,t)-path with backedges such that*
1. *for every vertex v of G_0, the number of vertices whose distance in G_0 from v is at most m is less than n^{ε};*
2. *there is no cycle of length less than m in G_0; and*
3. *however the spoiler colors G_0 with the c colors, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that each vertex has the same (d,q)-color in G_0 as in $G_0 - e$.*

Proof. By increasing m if necessary, we can assume without loss of generality that $m \geqslant q + 1$. We can also assume without loss of generality that $\varepsilon < \frac{1}{2}$. Take G_0 as in Theorem 7.1, and obtain G_0' by deleting some backedge on each short cycle (cycle of length less than m). Let $\varepsilon' = 2\varepsilon$. We shall let G_0' play the role of G_0 in the statement of the theorem we are now proving, and ε' play the role of ε. As we mentioned earlier, no point whose distance from a vertex is more than q affects the (d,q)-color of the vertex. It follows easily that if neither endpoint of a forward edge e is within distance q of an endpoint of a deleted backedge, then the (d,q)-color of every point in G_0' is the same as the (d,q)-color of the corresponding point in $G_0' - e$. So the theorem is proven if we show that the probability that some endpoint of a randomly selected forward edge e is within distance q of an endpoint of a deleted backedge is less than ε, when n is sufficiently large. Now the endpoints of deleted backedges each lie on short cycles. But the number of vertices that are within distance $m \geqslant q + 1$ of the short cycles is less than $(n^{\varepsilon})(n^{\varepsilon}) = n^{2\varepsilon}$, where the first factor of n^{ε} is an upper bound on the number of vertices w on short cycles, and the second factor of n^{ε} is an upper bound on the number of vertices within distance m of such a vertex w. The number of forward edges in G_0' is $n - 1$, since the number of vertices is n. So the probability that some endpoint of a randomly selected forward edge e is within distance q of an endpoint of a deleted backedge is at most $n^{2\varepsilon}/(n-1)$, which (since $\varepsilon < \frac{1}{2}$) is less than ε if n is sufficiently large. This was to be shown. $\square$

We can now give our simpler proof of the following theorem of Ajtai and Fagin [1].

Theorem 7.3. *Directed reachability is not in monadic* NP.

Proof. Let c (the number of colors) and r (the number of rounds) be given. Let $\mathscr{S}$ be the class of directed graphs that are (s,t)-connected. By Theorem 5.4, we need only show that the duplicator has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$. Let k be as in Theorem 4.1, and let ε be $1/(2k)$. Define f, d, q and m to be $\lfloor n^{\varepsilon} \rfloor$, r, k, and k, respectively. Let n be sufficiently large (we shall discuss how large later). As his first move in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$, the duplicator selects the graph G_0 guaranteed by Theorem 7.2, and the spoiler then colors G_0 with the c colors. Denote the resulting colored graph by $\mathbf{A}_0$. Let C be the total number of possible (d,q)-colors. Since C depends only on d and q, and since d and q depend only on r (since k depends only on r), it follows that C depends only on r. Let us call a forward edge e of $\mathbf{A}_0$ *good* if each vertex in $\mathbf{A}_0$ has the same (d,q)-color in $\mathbf{A}_0$ as in $\mathbf{A}_0 - e$. Since there are $n-1$ forward edges e, it follows from Theorem 7.2 that at least $(1-\varepsilon)(n-1)$ edges e are good. Since there are only C possible (d,q)-colors, there is some set S of at least $(1-\varepsilon)(n-1)/C$ good edges that all have the same color. This number $(1-\varepsilon)(n-1)/C$ is greater than $n^{1/2}+1 \geqslant f^k+1$ if n is sufficiently large (the last inequality holds since $f = \lfloor n^{\varepsilon} \rfloor \leqslant n^{1/(2k)}$). Select e to be a member of S. The duplicator now selects $G_0 - e$ as a member of $\overline{\mathscr{S}}$, and colors $G_0 - e$, vertex by vertex, with the same coloring as G_0. Denote the resulting colored graph $\mathbf{A}_0 - e$ by $\mathbf{A}_1$. The conditions of Theorem 4.1 are now satisfied, as we now show. The first three conditions of Theorem 4.1 follow immediately from the corresponding conditions of Theorem 7.2, by our choice of parameters. The fourth condition of Theorem 4.1 follows from our choice of e, and the fact that e is the only edge that is present in one graph and not in the other. Therefore, it follows from Theorem 4.1 and our choice of k that $\mathbf{A}_0 \sim_r \mathbf{A}_1$. Thus, the duplicator has a winning strategy in the remaining r-game. So the duplicator has a winning strategy in the Ajtai–Fagin (c,r)-game over $\mathscr{S}$, as desired. □

8. Allowing built-in relations

Ajtai and Fagin [1] proved that directed reachability is not in monadic NP, even in the presence of built-in relations from a large class, which includes the successor relation. Our techniques allow us to extend their results to a larger class of built-in relations. The significance of allowing built-in relations is discussed in the introduction. We now explain the definition.

For the purpose of considering built-in relations, it is convenient to restrict our attention to universes that are an initial prefix of the set of natural numbers. Thus, if the cardinality of the universe is n, then we assume that the universe is $\{0,\ldots,n-1\}$.

A particular collection of *built-in relations* is specified by an *auxiliary language* $\mathscr{L}'$ of relation symbols, along with, for each positive integer n, an interpretation of $\mathscr{L}'$ on the universe $\{0,\ldots,n-1\}$. Intuitively, each universe has associated with it a set of auxiliary relations. We denote the corresponding $\mathscr{L}'$-structure by $\Gamma(n)$. Intuitively, $\Gamma(n)$ is the structure composed of the built-in relations.

We now explain Ajtai and Fagin's result, which says that directed connectivity is not in monadic NP, even in the presence of certain families of built-in relations. Intuitively, they allow binary built-in relations with no short cycles and where every vertex has small degree.

Theorem 8.1 (Ajtai and Fagin [1]). *Assume that the built-in relations are all binary, and that $\xi(n) \to \infty$ and $\sigma(n) \to 0$ as $n \to \infty$. Assume also that $\Gamma(n)$ contains no cycle of length less than $\xi(n)$, and the degree of each point in $\Gamma(n)$ is at most $n^{\sigma(n)}$. Then directed reachability is not in monadic* NP, *even in the presence of the built-in relations.*

What does this result say? Let G be a graph (which we recall is defined to be a structure over the language $\mathscr{L}$ consisting of a single binary relation symbol and two constant symbols, s and t), with universe $\{0,\ldots,n-1\}$. By the *expansion* of G, we mean the structure $\widehat{G}$ over the expanded language $\mathscr{L} \cup \mathscr{L}'$, where the interpretation in $\widehat{G}$ of $\mathscr{L}$ is as in G, and the interpretation in $\widehat{G}$ of $\mathscr{L}'$ is as in $\Gamma(n)$. Theorem 8.1 says that as long as the built-in relations are restricted as described above, then there is no monadic Σ^1_1 sentence φ over the expanded language such that a graph G is (s,t)-connected iff $\widehat{G}$ satisfies φ.

We strengthen Theorem 8.1 by removing the restriction that the built-in relations be binary. Furthermore, we replace the assumption that $\Gamma(n)$ contains no short cycle (cycle of length less than $\xi(n)$) by an assumption that not very many points (at most $n^{\sigma(n)}$ points) lie on short cycles.

Theorem 8.2. *Assume that $\xi(n) \to \infty$ and $\sigma(n) \to 0$ as $n \to \infty$. Assume also that the number of points in $\Gamma(n)$ that lie on cycles of length less than $\xi(n)$ is at most $n^{\sigma(n)}$, and the degree of each point in $\Gamma(n)$ is at most $n^{\sigma(n)}$. Then directed reachability is not in monadic* NP, *even in the presence of the built-in relations.*

We now explain how to prove Theorem 8.2. Ajtai and Fagin give a version of Theorem 7.1 that applies to the case of built-in relations (under their assumptions on built-in relations). It is fairly straightforward to modify their proof to obtain the same theorem (Theorem 8.3 below), under our assumptions on built-in relations. When we talk about the (d,q)-color of a point in $\widehat{G_0}$, we take into consideration not only G_0, but also the built-in relations. By $\widehat{G_0 - e}$, we mean $\widehat{G}$, where G is $G_0 - e$. Thus, $\widehat{G_0 - e}$ is the result of deleting the edge e from the graph, but leaving the built-in relations the same.

Theorem 8.3. *Assume that c, d, q, and m are positive integers, and $\varepsilon > 0$. If n is sufficiently large, then there is a graph G_0 with universe $\{0,\ldots,n-1\}$ that is an (s,t)-path with backedges such that*

1. *for every vertex v of $\widehat{G_0}$, the number of vertices whose distance in $\widehat{G_0}$ from v is at most m is less than n^{ε};*
2. *the number of vertices that are on some cycle of size less than m in $\widehat{G_0}$ is less than n^{ε}; and*
3. *however the spoiler colors $\widehat{G_0}$ with the c colors, the probability is at least $1-\varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that each vertex has the same (d,q)-color in $\widehat{G_0}$ as in $\widehat{G_0}-e$.*

Just as in the case we considered earlier where there are no built-in relations, we prove Theorem 8.2 by giving a sufficient condition for the duplicator to have a winning strategy. Since we are now allowing small cycles, we must weaken the assumptions in Theorem 4.1. Let us define a point to be *k-isolated* if it is not within distance k of a cycle of length less than k. Intuitively, a point is k-isolated if it is not near any small cycle. A tuple in one of the relations in a structure is said to be k-isolated if each vertex in the tuple is k-isolated.

Theorem 8.4. *Let r, f be positive integers. There is a positive integer k that depends only on r such that $\mathbf{A}_0 \sim_r \mathbf{A}_1$ whenever $\mathbf{A}_0$ and $\mathbf{A}_1$ are structures of the same similarity type that have the same set of vertices and that satisfy the following conditions:*

1. *the degree of every vertex in $\mathbf{A}_0$ or $\mathbf{A}_1$ is at most f;*
2. *each vertex has the same (r,k)-color in $\mathbf{A}_0$ as in $\mathbf{A}_1$; and*
3. *if e is a tuple that is present in some relation in one structure but not in the corresponding relation in the other structure, then e is k-isolated, and there are at least f^k tuples in both of these relations that are k-isolated and have the same (r,k)-color as e.*

Proof. The proof is similar to that of Theorem 4.1, and the value of k is again 3^{2r}. We continue to assume for simplicity that $\mathbf{A}_0$ and $\mathbf{A}_1$ are colored graphs, rather than structures of arbitrary similarity type.

For an integer ℓ, let us say that a vertex i is *ℓ-safe* if in both graphs, all the tuples within a distance ℓ of this vertex are also present in the other graph. We observe that if the (r,ℓ)-color of a vertex is different in the two graphs, then it is not ℓ-safe. This in turn implies that there are "many" vertices in both graphs that have these two colors.

As before, the duplicator maintains an isomorphism invariant and a multiplicity invariant. The isomorphism invariant is the same as in the proof of Theorem 4.1. The multiplicity invariant, which we call the *strong multiplicity invariant*, is given below. The duplicator's strategy is to keep picking the same vertex in $\mathbf{A}_0$ (resp., $\mathbf{A}_1$) as the spoiler does in $\mathbf{A}_1$ (resp., $\mathbf{A}_0$) until the isomorphism invariant is threatened. Whenever the isomorphism invariant is threatened, the strong multiplicity invariant comes to the rescue, just as in the proof of Theorem 4.1.

- *Strong multiplicity invariant*: If $j \leqslant i$ is such that $p_j \neq q_j$, then both p_j and q_j are k_i-isolated, and there are at least f^{k_i} vertices that are k_i-isolated in both structures and that have the same $(r, 3^{2r-i})$-color as p_j and q_j, where $k_0 = 3^{r+1}$, and $k_i = k_{i-1} - 3^{r-i}$. (It is straightforward to verify that $k_i > 3^r$ for $i \leqslant r$.)

Assuming that the invariants have been maintained for $i < r$ rounds, we show how to maintain them for the next round. As in the proof of Theorem 4.1, we need to analyze many cases, and we give details only in the cases that require a new argument.

Suppose the spoiler picks a vertex p_{i+1} in $\mathbf{A}_0$ in round $i+1$ whose distance is more than 3^{r-i-1} from every previously chosen vertex in $\mathbf{A}_0$. If picking the corresponding vertex in $\mathbf{A}_1$ does not violate the isomorphism invariant, the duplicator picks it; this clearly maintains the strong multiplicity invariant. Assume therefore that the isomorphism invariant is violated. Since by assumption each vertex has the same (r,k)-color in both graphs, it follows that the corresponding vertex is close to (that is, within distance 3^{r-i-1} of) some other vertex q_j already picked. We now show how to find q_{i+1}. There are two cases, depending on whether or not $p_j = q_j$. If $p_j = q_j$, then since p_{i+1} is within distance 3^{r-i-1} of p_j in $\mathbf{A}_1$ but not in $\mathbf{A}_0$, it follows that p_{i+1} is not 3^{r-i-1}-safe. Let e be the edge that makes p_{i+1} unsafe. According to the third hypothesis of the theorem, e is k-isolated, so p_{i+1} is $(k-3^{r-i-1})$-isolated, and hence k_i-isolated. Also according to the third hypothesis of the theorem, there are f^k edges of the same (r,k) color as e in $\mathbf{A}_1$ that are k-isolated. Arguing as in part 2 of Fact 3 in the proof of Theorem 4.1, we see that since vertex p_{i+1} has distance at most 3^{r-i-1} from e, there are at least $f^k/f^{3^{r-i-1}} \geqslant f^{3^{2r}/2} \geqslant f^{k_{i+1}}$ vertices in $\mathbf{A}_0$ with the same $(r, k-3^{r-i-1})$ color (and hence the same $(r, 3^{2r-i-1})$-color) as p_{i+1} that are $(k-3^{r-i-1})$-isolated (and hence k_{i+1}-isolated). Since $|\bigcup_{j \leqslant i} \mathrm{Ball}(q_j, 3^{r-i-1})| \leqslant i f^{3^{r-i-1}} < f^{3^{2r}/2}$, at least one of these vertices is not within a distance 3^{r-i-1} of any of $q_1, \ldots, q_i$. The duplicator picks that vertex as q_{i+1}. It follows immediately from what we have shown that this maintains the strong multiplicity invariant (the isomorphism invariant is trivially maintained). Now assume $p_j \neq q_j$. By the strong multiplicity invariant, q_j is k_j-isolated, and so p_{i+1} is $(k_j - 3^{r-i-1})$-isolated, and hence k_{i+1}-isolated. Also by the strong multiplicity invariant, there are at least f^{k_j} vertices in $\mathbf{A}_1$ with the same $(r, 3^{2r-j})$-color as q_j that are all k_j-isolated. Arguing as in part 2 of Fact 3 in the proof of Theorem 4.1, we see that there are at least $f^{k_j}/f^{3^{r-i-1}} \geqslant f^{k_{i+1}}$ vertices in $\mathbf{A}_1$ with the same $(r, 3^{2r-j} - 3^{r-i-1})$-color (and hence the same $(r, 3^{2r-i-1})$-color) as p_{i+1} that are $(k_j - 3^{r-i-1})$-isolated, and therefore k_{i+1}-isolated. Since $|\bigcup_{j \leqslant i} \mathrm{Ball}(q_j, 3^{r-i-1})| \leqslant i f^{3^{r-i-1}} < r^{3^r} < f^{k_{i+1}}$, one of these vertices is not within a distance 3^{r-i-1} of any of $q_1, \ldots, q_i$. The duplicator picks that vertex as q_{i+1}, thus maintaining the strong multiplicity invariant.

Now suppose that the spoiler has picked a vertex in $\mathrm{Ball}(p_j, 3^{r-i-1})$ for some $j \leqslant i$. Assume first that $p_j = q_j$. If p_j, q_j are k_j-safe, then the duplicator trivially maintains the isomorphism invariant by picking for q_{i+1} the vertex in $\mathbf{A}_1$ that corresponds to p_{i+1}; this is because $k_j > 3^r$, so $\mathrm{Ball}(p_j, 3^r)$ and $\mathrm{Ball}(q_j, 3^r)$ look exactly the same. On the other hand, if p_j, q_j are not k_j-safe, then their neighborhoods are tree-like, and techniques similar to those in the proof of Theorem 4.1 let us maintain the invariants.

This leaves the case that $p_j \neq q_j$. By the strong multiplicity invariant, p_j and q_j are k_j-isolated, and hence their immediate neighborhoods are tree-like. Then the duplicator can follow the strategy in the proof of Theorem 4.1 to pick a vertex q_{i+1} in $\mathrm{Ball}(q_j, 3^{r-i-1})$ of the same $(d, 3^{2r-i} - 3^{r-i-1})$-color (and hence the same $(d, 3^{2r-i-1})$-color) as p_{i+1}. Now we show that this maintains the strong multiplicity invariant (the isomorphism invariant is clearly maintained). Note that (i) p_{i+1} and q_{i+1} are $(k_j - 3^{r-i-1})$-isolated (and hence k_{i+1}-isolated), since p_j and q_j are k_j-isolated, and (ii) by part 2 of Fact 3, the multiplicity of p_{i+1} (resp., q_{i+1}) is at least $1/f^{3^{r-i-1}}$ times the multiplicity of p_j (resp., q_j). Clearly, (i) and (ii) imply that the strong multiplicity invariant is maintained. □

We can think of Theorem 8.4 as a modification of Theorem 4.1 where we allow small cycles, as long as they are not near edges that appear in one structure but not the other.

The proof of Theorem 8.2 now follows essentially the same outline as the proof of Theorem 7.3, except that we use Theorems 8.3 and 8.4 instead of Theorems 7.2 and 4.1. Furthermore, instead of using Theorem 5.4, we use the natural variation of Theorem 5.4, that holds in the case of built-in relations. The only essential change in the proof is that we take advantage of the fact that almost all forward edges of G_0 in Theorem 8.3 are m-isolated (this follows immediately from parts (1) and (2) of Theorem 8.3). Thus in addition to all the other properties, $1-\varepsilon$ fraction of the forward edges are also m-isolated. We then modify the proof of Theorem 7.3 by taking the forward edge e to be m-isolated, in addition to the other properties demanded of it. The straightforward details are left to the reader.

9. Summary

We present a strong new condition that guarantees that the duplicator has a winning strategy in an Ehrenfeucht–Fraïssé game. This gives a greatly simplified proof that directed reachability is not in monadic NP. Although our condition was designed for the directed reachability question, its versatility is shown by the fact that it can also be used for proving that connectivity is not in monadic NP (in a proof very different from that used for directed reachability). Furthermore, a slight variation of our condition (where we allow small cycles, as long as they are not near any edges that appear in one structure but not the other) leads to new, strengthened results on descriptive complexity in the presence of built-in relations.

Acknowledgements

The authors are grateful to Joe Halpern, Phokion Kolaitis, Larry Stockmeyer, and Moshe Vardi for giving us useful comments on an earlier version of the paper. Furthermore, we thank Yuri Gurevich for suggesting that we define the winning condition

for the duplicator on structures over arbitrary languages, rather than restricting our attention to graphs.

References

[1] M. Ajtai and R. Fagin, Reachability is harder for directed than for undirected finite graphs, *J. Symbolic Logic* **55** (1990) 113–150.

[2] C. Beeri, P.C. Kanellakis, F. Bancilhon and R. Ramakrishnan, Bounds on the propagation of selection into logic programs, in: *Proc. 6th ACM Symp. on Principles of Database Systems* (1987) 214–226.

[3] C. Berge, *Graphs and Hypergraphs* (North-Holland, New York, 1976).

[4] A. Chandra and D. Harel, Structure and complexity of relational queries, *J. Comput. System Sci.* **25** (1982) 99–128.

[5] S. Cosmadakis, Logical reducibility and monadic NP, in: *Proc. 34th IEEE Symp. on Foundations of Computer Science* (1993) 52–61.

[6] M. de Rougemont, Second-order and inductive definability on finite structures, *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik* **33** (1987) 47–63.

[7] A. Ehrenfeucht, An application of games to the completeness problem for formalized theories, *Fund. Math.* **49** (1961) 129–141.

[8] H.B. Enderton, *A Mathematical Introduction to Logic* (Academic Press, New York, 1972).

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

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

[11] R. Fagin, Degrees of acyclicity for hypergraphs and relational detabase schemes, *J. ACM* **30** (1983) 514–550.

[12] R. Fagin, L. Stockmeyer and M.Y. Vardi, On monadic NP vs. monadic co-NP, *Inform. and Comput.* **120**(1) (1995) 78–92.

[13] R. Fraïssé, Sur quelques classifications des systèmes de relations, *Publ. Sci. Univ. Alger. Sér. A* 1 (1954) 35–182.

[14] H. Gaifman, On local and nonlocal properties, in: J. Sterm, ed., *Logic Colloquium'81* (North Holland, Amstendam, 1982) 105–135.

[15] M.R. Garey and D.S. Johnson, *Computers and Intractability – A Guide to the Theory of NP-Completeness* (W.H. Freeman, New York, 1979).

[16] W. Hanf, Model-theoretic methods in the study of elementary logic, in: J. Addison, L. Henkin and A. Tarski, eds., *The Theory of Models* (North Holland, Amstendan, 1965) 132–145.

[17] N. Immerman, Relational queries computable in polynomial time, *Inform. and Control* **68** (1986) 86–104.

[18] N. Immerman, Descriptive and computational complexity, in: J. Hartmanis, ed., *Computational Complexity Theory, Proc. Symp. Applied Math., Vol 38* (American Mathematical Society, Providence, RI 1989) 75–91.

[19] T. Schwentick, Graph connjectivity and monadic NP, in: *Proc. 35th IEEE Symp. on Foundations of Computer Science* (1994) 614–622.

[20] J.R. Shoenfield, *Mathematical Logic* (Addison-Wesley, Reading MA, 1967).

[21] G. Turań, On the definability of properties of finite graphs, *Discrete Math.* **49** (1984) 291–302.

[22] M.Y. Vardi, The complexity of relational query languages, in: *Proc. 14th ACM Symp. on Theory of Computing* (1982) 137–146.