On Monadic NP vs Monadic co-NP*

RONALD FAGIN, LARRY J. STOCKMEYER, AND MOSHE Y. VARDI[†]

IBM Research Division, Almaden Research Center, 650 Harry Road, San Jose, California 95120-6099 E-mail: {fagin, stock}@almaden.ibm.com, vardi@cs.rice.edu

It is a well-known result of Fagin that the complexity class NP coincides with the class of problems expressible in existential second-order logic (Σ_1^1) . Monadic NP is the class of problems expressible in monadic Σ_1^1 , i.e., Σ_1^1 with the restriction that the second-order quantifiers range only over sets (as opposed to ranging over, say, binary relations). We prove that connectivity of finite graphs is not in monadic NP, even in the presence of arbitrary built-in relations of moderate degree (that is, degree $(\log n)^{o(1)}$). This extends earlier results of Fagin and de Rougemont. Our proof uses a combination of three techniques: (1) an old technique of Hanf for showing that two (infinite) structures agree on all first-order sentences, under certain conditions, (2) a recent new approach to second-order Ehrenfeucht-Fraissé games by Ajtai and Fagin, and (3) playing Ehrenfeucht-Fraissé games over random structures (this was also used by Ajtai and Fagin). Regarding (1), we give a version of Hanf's result that is better suited for use as a tool in inexpressibility proofs for classes of finite structures. The power of these techniques is further demonstrated by using them (actually, using just the first two techniques) to give a very simple proof of the separation of monadic NP from monadic co-NP without the presence of built-in relations. © 1995 Academic Press, Inc.

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. Complexity theory traditionally has focused on the computational complexity of problems. A more recent branch of complexity theory focuses on the *descriptive complexity* of problems, which is the complexity of describing problems in some logical formalism [Imm89]. One of the exciting developments in complexity theory is the discovery of a very intimate connection between computational and descriptive complexity.

This intimate connection was first discovered by Fagin, who showed [Fag74] (cf. [JS74]) 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 . Stockmeyer then observed that this could be extended to give a tight correspondence between the polynomial-time hierarchy and second-order logic [Sto77].

* A condensed version of this paper appeared in "Proceedings of the 8th IEEE Conference on Structure in Complexity Theory," 1993, pp. 19–30.

The next discovery was by Immerman and Vardi, who proved that the complexity class P coincides with the class of properties of finite ordered structures expressible in fixpoint logic [Imm86, Var82]. The connection between descriptive and computational complexity, typically referred to as the connection between "logic and complexity", was then proclaimed by Immerman [Imm87], and studied by many researchers; see [Imm89] for a survey. This connection is considered to be one of the major features of finite-model theory; see [Fag93].

A consequence of the connection between NP and existential second-order logic is that 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 $\Sigma_1^1 = \Pi_1^1$. This equivalence of questions in computational and descriptive complexity is one of the major features of the connection between the two branches of complexity theory. It holds the promise that techniques from one domain could be brought to bear on questions in the other domain. In particular, there is a standard technique in finite-model theory for proving separation results: Ehrenfeucht-Fraïssé games. It is known that $\Sigma_1^1 \neq \Pi_1^1$ if and only if such a separation can be proven via second-order Ehrenfeucht-Fraïssé games [Fag75a]. Unfortunately, "playing" second-order Ehrenfeucht-Fraïssé games is very difficult, and the above promise is essentially still largely unfulfilled; for example, the equivalence between the NP = co-NP question and the $\Sigma_1^1 = \Pi_1^1$ question has not so far led to any progress on either of these questions.

One way of attacking these difficult questions is to restrict the classes under consideration. Instead of considering Σ_1^1 (=NP) and $\Pi_1^1(=co-NP)$ in their full generality, we could consider the monadic restriction of these classes, i.e., the restriction obtained by allowing second-order quantification only over sets (as opposed to quantification over, say, binary relations). We refer to the restricted classes as monadic Σ_1^1 or monadic NP (resp., monadic Π_1^1 or monadic co-NP). (It should be noted that, in spite of its severely restricted syntax, monadic NP does contain NP-complete problems, such as 3-colorability and satisfiability.) The hope is that the restriction to the monadic classes will yield more tractable questions and will serve as a training ground for attacking the problems in their full generality.

[†] Current address: Department of Computer Science, Rice University, P.O. Box 1892, Houston, TX 77251-1892.

This line of attack was pursued by Fagin in [Fag75a], where he separated monadic NP from monadic co-NP. Specifically, he showed that connectivity of finite graphs is not in monadic NP, although it is easy to see that it is in monadic co-NP. This result was the first lower bound in descriptive-complexity theory. It was also the first significant demonstration of the weakness of first-order logic over finite structures, since it implies that connectivity of finite graphs is not expressible in first-order logic (the inexpressibility of connectivity of general graphs in first-order logic is a trivial consequence of the Compactness Theorem). This consequence was rediscovered later by Aho and Ullman [AU79] and inspired a great deal of research in the theory of database queries (cf. [Cha88]) and in finite-model theory.

To separate monadic NP from monadic co-NP, Fagin extended the theory of Ehrenfeucht-Fraïssé games to monadic Σ_1^1 . In the standard Ehrenfeucht-Fraïssé game over a pair **A**, **B** of structures, two players, the spoiler and the duplicator, take turns placing pebbles on elements of the structures.¹ In the game for monadic Σ_1^1 , the spoiler starts by coloring the elements of **A**, the duplicator responds by coloring the elements of **B**, and the two players then follow the standard game. To show that connectivity of finite graphs is not expressible in monadic Σ_1^1 , Fagin used the generalized game over a pair **A**, **B** of graphs, where **A** consists of a single cycle and **B** consists of two cycles. The separation of monadic NP and monadic co-NP now follows, since, as we noted above, connectivity is in monadic co-NP.

One essential difference between NP and monadic NP is that in NP one can assume the existence of certain built-in relations on the domain, such as successor or linear order, since the existence of such relations can be expressed by a second-order existential quantifier. This is not the case for monadic NP, which is one of the reasons for the weakness of this class. 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.

Extending the techniques to handle built-in relations is important, since some connections between computational complexity and descriptive complexity are known to hold only if there is a built-in successor relation (or a built-in linear order). For example, as we noted 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 [CH82].

There is another reason (besides our interest in successor relations) to allow built-in relations. Proving that a problem is not in monadic NP shows that the problem cannot be captured in a certain *uniform* way, where we think of a fixed monadic Σ_1^1 sentence as a uniform description. 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). So allowing builtin relations makes nonexpressibility results that much more powerful. We note that first-order logic, in the presence of arbitrary built-in relations, is precisely (nonuniform) AC⁰, that is, properties that can be recognized by a family of polynomial-size circuits with bounded depth [Imm87]. It follows, for example, that the graph property "the number of edges is even" cannot be expressed in first-order logic with arbitrary built-in relations, since this property is not in AC⁰ [Ajt83, FSS84].

Unfortunately, extending Fagin's result, that connectivity is not in monadic NP, to allow (certain) built-in relations is not easy. The hard part in Fagin's proof is showing that the duplicator has a winning strategy. There are several parts to the duplicator's winning strategy: his coloring strategy, his pebbling strategy, and (depending on the version of the game we consider) also his strategy in the choice of graphs to play the game over (the graphs are simply disjoint unions of cycles, but the size of the cycles is an issue). All of these parts of the duplicator's winning strategy in Fagin's proof are very complicated. The complexity of the proof makes it quite hard to extend it to built-in relations. Such an extension was accomplished by de Rougemont, who proved that connectivity is not in monadic NP with a built-in successor relation [dR87] (by considering graphs that are substantially more complicated than the cycles in Fagin's proof).

Our goal in this paper is to provide new tools for separation proofs. This provides us not only with a simple and elegant proof of Fagin's result, but also with its extension to arbitrary built-in relations of moderate degree (that is, degree $(\log n)^{o(1)}$). Such built-in relations include successor relations, but not linear orders.² We accomplish this by using three tools: an old but relatively unknown technique by Hanf [Han65] for showing that the duplicator has a winning strategy in certain situations, a recent new approach to Ehrenfeucht-Fraïssé games by Ajtai and Fagin [AF90], and the idea (used also by Ajtai and Fagin) of having the duplicator select structures at random.

The basic idea in the approach of Ajtai and Fagin is not to view the pair **A**, **B** of structures as an input to the game. Rather, they should be viewed as selected by the duplicator. According to this view, to prove that a certain property P is not expressible by a monadic Σ_1^1 sentence, the game proceeds as follows. The duplicator selects a pair **A**, **B** of structures such that P holds for **A** and fails for **B**. The two players then play the generalized game over **A**, **B**. 2

¹ Following Joel Spencer [Spe91], we shall refer to the two players in an Ehrenfeucht-Fraïssé game as "the spoiler" and "the duplicator," rather than the more usual but less suggestive "player I" and "player II."

 $^{^{2}}$ Note added in proof. Schwentick [Sch94a, Sch94b] has recently shown that connectivity is not in monadic NP, even in the presence of a linear order.

However, once the selection of the structures is viewed as a move in the game, it is quite natural to consider interleaving this move with the other moves. Ajtai and Fagin considered the following interleaving. The duplicator first selects the structure **A**, which is then colored by the spoiler. The duplicator then selects the structure **B** and colors it. The two players then play the standard game over the colored structures **A**, **B**. Note that this variant handicaps the spoiler and helps the duplicator. Nevertheless, Ajtai and Fagin showed that it suffices to consider their variant when trying to prove lower bounds on expressibility for monadic NP. The advantage of using this variant of the game is that it is tilted in favor of the duplicator, and therefore it is much easier to describe a winning strategy for the duplicator.

Ajtai and Fagin introduced another powerful idea that we use: having the duplicator select structures *at random*; it suffices to show that the probability of winning is nonzero. (Actually, both here and in Ajtai and Fagin's proof, it is shown that the probability of winning is not just nonzero, but nearly one.)

As we will show, in the case of connectivity, the Ajtai-Fagin game makes the coloring step for the duplicator easy: the duplicator can essentially "copy" the coloring of the spoiler. This leaves the other difficult part of the strategy—how the duplicator responds to pebble moves by the spoiler. In our proof, we use a "library subroutine" based on Hanf's technique that gives the duplicator's winning strategy for responding to pebble moves by the spoiler.

We note that Aitai [Ait83] previously proved separation between monadic NP and monadic co-NP allowing arbitrary built-in relations. In fact, Aitai proved the very strong separation result that there is a (somewhat artificial) property of graphs, which belongs to monadic co-NP (with no built-in relations), but which does not belong to monadic NP even in the presence of arbitrary built-in relations. Ajtai and Fagin [AF90] proved a separation involving (s, t)-connectivity (otherwise known as directed reachability): they showed that although this problem is in monadic co-NP, it is not in monadic NP, even in the presence of binary built-in relations of degree $n^{o(1)}$, as long as these built-in relations have no "small cycles." In our separation result, there is no restriction on the length of cycles or on the arity of the built-in relations, but we can only allow degree $(\log n)^{o(1)}$.

Recently, Arora and Fagin [AF94] found another technique, different from Hanf's, for showing that the duplicator has a winning strategy in certain situations. They showed the usefulness of this new tool in a way parallel to ours. Specifically, they used this technique in two ways: (1) they gave a proof that directed (s, t)-connectivity is not in monadic NP that is much easier than the earlier proof by Ajtai and Fagin, and (2) they showed that directed (s, t)connectivity is not in monadic NP in the presence of a larger class of built-in relations than was known before. In (2), they allow built-in relations of arbitrary arity and they allow small cycles, as long as not too many vertices lie on the small cycles. They also showed that they can replace Hanf's technique by their technique, in our proof of Fagin's result that connectivity is not in monadic NP.

Turán [Tur84] has taken Fagin's result in another direction by showing that connectivity is not expressible in existential monadic second-order logic if we can (existentially) quantify over sets of edges of G as well as sets of vertices of G (in Fagin's result, quantification only over sets of vertices is allowed). Essentially, this amounts to representing a graph as a set of vertices, a set of edges, and an incidence relation between vertices and edges. In contrast, we represent a graph as a set of vertices and an edge relation. The former representation of graphs is, in fact, the representation used by Courcelle in, for example, [Cou90], where quantification over both vertices and edges of the graph are allowed. We remark that our proofs that connectivity is not in monadic NP (both the simple proof not allowing built-in relations and the more complicated proof allowing built-in relations of moderate degree) still work with very minor modification for this alternate approach.

We do not know whether our restriction on the built-in relations (that they be of moderate degree) is essential: we consider it possible that connectivity is not in monadic NP, even in the presence of arbitrary built-in relations of arbitrary degree and arity (sometimes called "a polynomial amount of advice").

2. DEFINITIONS AND CONVENTIONS

A language \mathcal{L} (sometimes called a similarity type, a signature, or a vocabulary) is a finite set $\{P_1, ..., P_s\}$ of relation symbols, each of which has an arity.

An \mathcal{L} -structure (or structure over \mathcal{L} , or simply structure) is a set A (called the *universe*), along with a mapping associating a relation R_i over A with each $P_i \in \mathcal{L}$, where R_i has the same arity as P_i , for $1 \leq i \leq s$. We may call R_i the *interpretation of* P_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. We note that all of our results hold whether or not we restrict our attention to finite structures. (The fact that connectivity is not in Σ_1^1 , monadic or otherwise, even in the presence of a built-in linear order, has an extremely simple proof using the Compactness Theorem in the case where infinite structures are allowed.)

In this paper, we are especially interested in *graphs* and *colored graphs*. Graphs are simply structures where the language consists of a single binary relation symbol. Although such structures are in general directed graphs, we often view the structure as an undirected graph by ignoring the directions of the edges. Colored graphs are structures where the language consists of a single binary relation

symbol and some number of unary relation symbols. If G is a colored graph, where the interpretations of the unary relation symbols in the language are $U_1, ..., U_k$, 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, intuitively, there are 2^k possible colors.

For definitions of a first-order sentence (where, intuitively, the only quantification is over members of the universe, and not over, say, sets of members of the universe), and what it means for a structure **A** to *satisfy* a sentence σ , written $\mathbf{A} \models \sigma$, see Enderton [End72] or Shoenfield [Sho67]. We note that equality is treated as a special relation symbol, which is not considered to be a member of the language \mathcal{L} , and which always has the standard interpretation.

When we pass from first-order logic to second-order logic, we allow quantification over sets and relations. In particular, a Σ_1^1 sentence is a sentence of the form $\exists A_1 \cdots \exists A_k \psi$, where ψ is first-order and where the A_i 's are relation symbols. As an example, we now construct a Σ_1^1 sentence that says that a graph (with edge relation denoted by E) is 3-colorable. Let E'xy denote $Exy \lor Eyx$. 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

$$\forall x((A_1x \land \neg A_2x \land \neg A_3x) \lor (\neg A_1x \land A_2x \land \neg A_3x) \\ \lor (\neg A_1x \land \neg A_2x \land A_3x)).$$

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

$$\forall x \; \forall y ((A_1 x \land A_1 y \Rightarrow \neg E' xy) \land (A_2 x \land A_2 y \Rightarrow \neg E' xy) \land (A_3 x \land A_3 y \Rightarrow \neg E' xy)).$$

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

As another example, which is very relevant for this paper, we now show that the class of graphs that are not connected is Σ_1^1 (this demonstration is from [Fag75a]). Let ψ_1 say "The set A is nonempty and its complement is nonempty," that is, $\exists x \exists y(Ax \land \neg Ay)$. Let ψ_2 say "There is no edge between A and its complement," that is, $\forall x \forall y((Ax \land \neg Ay))$ $\Rightarrow \neg E'xy$). It is clear that the Σ_1^1 sentence $\exists A(\psi_1 \land \psi_2)$ characterizes the class of graphs that are not connected.

A Σ_1^1 sentence $\exists A_1 \cdots \exists A_k \psi$, 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{G} of \mathscr{L} -structures is said to be (*monadic*) Σ_1^1 if it is the class of all \mathscr{L} -structures that obey some fixed (monadic) Σ_1^1 sentence. A (monadic) Σ_1^1 class is also called a (monadic) generalized spectrum. One reason that Σ_1^1 classes are of great interest is the result [Fag74] that the collection of Σ_1^1 classes coincides with the complexity class NP. For this reason, we refer to the collection of monadic Σ_1^1 classes as *monadic NP*. We often refer to a class of graphs by a defining property, for example, 3-colorability or connectivity. As we saw above, 3-colorability and nonconnectivity are in monadic NP. Note that 3-colorability is an NP-complete property [GJ79]. Thus, monadic NP includes NP-complete properties. Let us define a class to be in monadic co-NP if its complement is in monadic NP. For example, since nonconnectivity is in monadic NP, it follows that connectivity is in monadic co-NP. This is of interest, because one result of this paper is a simple proof of Fagin's result that connectivity is not in monadic NP (and an extension of this result where we allow certain built-in relations). In particular, monadic NP and monadic co-NP are not the same.

3. EHRENFEUCHT-FRAÏSSÉ GAMES

Among the few tools of model theory that "survive" when we restrict our attention to finite structures are Ehrenfeucht– Fraïssé-type games [Ehr61, Fra54]. For an introduction to Ehrenfeucht–Fraïssé games and some of their applications to finite-model theory, see [AF90, 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. For ease in description, we shall restrict our attention to colored graphs, but everything we say generalizes easily to arbitrary structures. There are two players, called the spoiler and the duplicator, and two colored graphs, G_0 and G_1 . In the first round, the spoiler selects a point in one of the two colored graphs, and the duplicator selects a point in the other colored graph. Let a_1 be the point selected in G_0 , and let b_1 be the point selected in G_1 . Then the second round begins, and again, the spoiler selects a point in one of the two colored graphs, and the duplicator selects a point in the other colored graph. Let a_2 be the point selected in G_0 , and let b_2 be the point selected in G_1 . This continues for r rounds. The duplicator wins if the colored subgraph of G_0 induced by $\langle a_1, ..., a_r \rangle$ is isomorphic to the colored subgraph of G_1 induced by $\langle b_1, ..., b_r \rangle$, under the function that maps a_i onto b_i for $1 \le i \le r$. That is, for the duplicator to win, (a) $a_i = a_i$ iff $b_i = b_i$, for each *i*, *j*; (b) (a_i, a_i) is an edge in G_0 iff (b_i, b_j) is an edge in G_1 , for each *i*, *j*; and (c) a_i has the same color as b_i , 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 $G_0 \sim_r G_1$. In this case, intuitively, G_0 and G_1 are indistinguishable by an r-game.

The following important theorem (from [Ehr61, Fra54]) shows why these games are of interest. If \mathscr{S} is a class of colored graphs, then let $\overline{\mathscr{S}}$ be the complement of \mathscr{S} , that is, the class of colored graphs not in \mathscr{S} .

THEOREM 3.1. \mathscr{S} is first-order definable iff there is r such that whenever $G_0 \in \mathscr{S}$ and $G_1 \in \overline{\mathscr{S}}$, then the spoiler has a winning strategy in the r-game over G_0 , G_1 .

We now discuss a more complicated game, which is a *c*-color, *r*-round, monadic NP game, and which we shall call a (c, r)-game for short. This game was introduced in [Fag75a] to prove that connectivity is not in monadic NP. We start with two graphs G_0 and G_1 (in this case, not colored). Let C be a set of c distinct colors. The spoiler first colors each of the points of G_0 , using the colors in C, and then the duplicator colors each of the points of G_1 , using the colors in C. 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 the colored subgraph of G_0 induced by $\langle a_1, ..., a_r \rangle$ is isomorphic to the colored subgraph of G_1 induced by $\langle b_1, ..., b_r \rangle$ (under the function that maps a_i onto b_i for $1 \leq i \leq r$).

The following theorem (from [Fag75a]) is analogous to Theorem 3.1.

THEOREM 3.2. A class \mathscr{S} of graphs is in monadic NP iff there are c, r such that whenever $G_0 \in \mathscr{S}$ and $G_1 \in \overline{\mathscr{S}}$, then the spoiler has a winning strategy in the (c, r)-game over G_0, G_1 .

In [Fag75a] it is shown that given c and r, there is a graph G_0 that is a cycle, and a graph G_1 that is the disjoint union of two cycles, such that the duplicator has a winning strategy in the (c, r)-game over G_0, G_1 . Since G_0 is connected and G_1 is not, it follows from Theorem 3.2 that connectivity is not in monadic NP.

In addition to considering games over pairs G_0 , G_1 of graphs, Ajtai and Fagin [AF90] found it convenient, for reasons we shall see shortly, to consider games over a class \mathscr{S} . The rules of an *r*-game over \mathscr{S} are as follows. The duplicator begins by selecting a member of \mathscr{S} to be G_0 , and a member of $\overline{\mathscr{S}}$ to be G_1 . The players then play an *r*-game over G_0 , G_1 to determine the winner. Similarly, we can define a (c, r)-game over \mathscr{S} . The rules are as follows.

- 1. The duplicator selects a member of \mathscr{S} to be G_0 .
- 2. The duplicator selects a member of $\overline{\mathcal{P}}$ to be G_1 .
- 3. The spoiler colors G_0 with the *c* colors.
- 4. The duplicator colors G_1 with the c colors.

5. The spoiler and duplicator play an r-game on the colored G_0, G_1 .

The next theorem follows easily from Theorems 3.1 and 3.2.

THEOREM 3.3. (a) \mathscr{S} is first-order definable iff there is r such that the spoiler has a winning strategy in the r-game over \mathscr{S} .

(b) \mathscr{S} is in monadic NP iff there are c, r such that the spoiler has a winning strategy in the (c, r)-game over \mathscr{S} .

We now explain why Ajtai and Fagin allow G_0 and G_1 to be selected by the duplicator, rather than inputs to the game. A (directed) graph with distinguished points s, t is said to be (s, t)-connected if there is a directed path in the graph from s to t. Ajtai and Fagin wished to prove that directed (s, t)-connectivity is not in monadic NP, but they did not see how to prove this by using (c, r)-games. By considering the choice of G_0 and G_1 to be moves of the duplicator, rather than inputs to the game, they were able to define a variation of (c, r)-games, in which the choice of G_1 by the duplicator is delayed until after the spoiler has colored G_0 . They successfully used the new game to prove the desired result (that directed (s, t)-connectivity is not in monadic NP). Their new game, which we call the Aitai-Fagin (c, r)-game, is, on the face of it, easier for the duplicator to win. The rules of the new game are obtained from the rules of the (c, r)-game by reversing the order of two of the moves. Specifically, the rules of the Ajtai-Fagin (c, r)-game are as follows.

- 1. The duplicator selects a member of \mathcal{S} to be G_0 .
- 2. The spoiler colors G_0 with the *c* colors.
- 3. The duplicator selects a member of $\overline{\mathscr{G}}$ to be G_1 .
- 4. The duplicator colors G_1 with the c colors.

5. The spoiler and duplicator play an *r*-game on the colored G_0, G_1 .

The winner is decided as before. Thus, in the Ajtai-Fagin (c, r)-game, the spoiler must commit himself to a coloring of G_0 with the *c* colors before knowing what G_1 is. In order to contrast it with the Ajtai-Fagin (c, r)-game, we may sometimes refer to the (c, r)-game as the *original* (c, r)-game. In spite of the fact that it seems to be harder for the spoiler to win the Ajtai-Fagin (c, r)-game than the original (c, r)-game, we have the following analogue [AF90] to Theorem 3.3(b).

THEOREM 3.4. \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} .

We will make use of Theorem 3.4 to give a simple proof that connectivity is not in monadic NP (and to extend to allowing certain built-in relations).

4. HANF'S TECHNIQUE

In this section, we shall provide a simple but very useful sufficient condition for guaranteeing that $A \sim_r B$ for two structures A, B. The proof is based on a technique of Hanf [Han65].

Let A be an \mathscr{L} -structure, where $\mathscr{L} = \{P_1, ..., P_s\}$, and where R_i is the interpretation in A of the relation symbol P_i , for $1 \le i \le s$. Let a and b be two points in (the universe of) A. We say that a and b are *adjacent* (in A) if either a = b or there is some R_i and some tuple t such that $t \in R_i$ and such that a and b are entries in the tuple t. Intuitively, two points a and b are adjacent if they are either identical or directly related by some relation of A. The *degree* of a point a is the cardinality of the set of points adjacent to a but not equal to a. By $A \upharpoonright X$ for a subset X of the universe of A, we mean the structure with universe X where the interpretation of P_i is the set of tuples t in R_i such that every entry of t is in X, for $1 \le i \le s$.

Essentially following Hanf, we define the *neighborhood* Nbd(d, a) of radius d about a recursively as follows:

 $Nbd(1, a) = \{a\}$

 $Nbd(d+1, a) = \{x \mid x \text{ is adjacent to some } b \in Nbd(d, a)\}.$

It is helpful to think of these neighborhoods as *open* spheres. Thus, intuitively, Nbd(d, a) consists of all points whose distance from a is *strictly* less than d. Note that because a is adjacent to itself, we have that $Nbd(d, a) \subseteq Nbd(d+1, a)$.

Also following Hanf, we define the *d-type* of a point *a* to be the isomorphism type of the neighborhood of radius *d* about *a* with *a* as a distinguished point. Thus, the points *a* in **A** and *b* in **B** have the same *d*-type precisely if $\mathbf{A} \upharpoonright Nbd(d, a) \cong \mathbf{B} \upharpoonright Nbd(d, b)$, under an isomorphism mapping *a* to *b*.

Let d, m be positive integers. We say that \mathcal{L} -structures A and B are (d, m)-equivalent if for every d-type τ , either A and B have the same number of points with d-type τ , or else both have at least m points with d-type τ . Intuitively, A and B are (d, m)-equivalent if for every d-type τ , they have the same number of points with d-type τ , where we can count only as high as m. We say that the structures A and B are d-equivalent if for every d-type τ , they have exactly the same number of points with d-type τ .

We need two simple lemmas.

LEMMA 4.1. If A and B are (d, m)-equivalent, and if d > d', then A and B are (d', m)-equivalent.

Proof. Let τ' be a d'-type. Let us say that a d-type τ refines τ' , and write $\tau > \tau'$, if every point with d-type τ also has d'-type τ' . Since d > d', it follows easily that for every point with d'-type τ' , there is some d-type τ that refines τ' . Define count(\mathbf{A}, τ) to be the number of points in \mathbf{A} with d-type τ (and similarly for count(\mathbf{B}, τ), count(\mathbf{A}, τ'), and count(\mathbf{B}, τ')). It follows easily from our remarks, and from the fact that every point has exactly one d-type, that

$$\operatorname{count}(\mathbf{A}, \tau') = \sum_{\tau > \tau'} \operatorname{count}(\mathbf{A}, \tau).$$
(1)

Identically,

$$\operatorname{count}(\mathbf{B}, \tau') = \sum_{\tau \succ \tau'} \operatorname{count}(\mathbf{B}, \tau).$$
(2)

Since A and B are (d, m)-equivalent, the right-hand sides (and hence the left-hand sides) of Eqs. (1) and (2) are either the same, or both at least m. Thus, count(A, τ') and count(B, τ') are either the same, or both at least m. Since τ' is an arbitrary d'-type, it follows that A and B are (d', m)equivalent.

LEMMA 4.2. Assume that $f \ge 2$. The size of a neighborhood of radius d in a structure where every point has degree at most f is less than f^d .

Proof. It is easy to see that the size of a neighborhood of radius d in a structure where every point has degree at most f is at most

$$\frac{1+f+f^2+\dots+f^{d-1}=(f^d-1)}{(f-1)< f^d}.$$

The next theorem (Theorem 4.3) is a key tool in our proof that connectivity is not in monadic NP (including the extension to allowing certain built-in relations). We give it in slightly more generality than we need, since we believe that it can be a useful tool in the future. The simpler version of the theorem that we actually use is then obtained as an immediate corollary.

THEOREM 4.3. Let r, f be positive integers. There are positive integers d, m, where d depends only on r, such that whenever **A** and **B** are (d, m)-equivalent structures where every point has degree at most f, then $\mathbf{A} \sim \mathbf{B}$.

Proof. We can assume without loss of generality that $f \ge 2$. We can also assume without loss of generality that the universes of **A** and **B** are disjoint. Let $d = 3^{r-1}$, and let $m = r \cdot f^{d-1}$. Assume that **A** and **B** are (d, m)-equivalent structures where every point has degree at most f. We now describe a winning strategy for the duplicator in an r-game over the structures **A**, **B**. The duplicator's strategy is to ensure that after j rounds, if $a_1, ..., a_j$ (resp. $b_1, ..., b_j$) are the points selected in **A** (resp. **B**), then a certain condition, which we call the *j*-matching condition holds:

j-matching condition: $\mathbf{A} \upharpoonright (\bigcup_{i \leq j} Nbd(3^{r-j}, a_i)) \cong \mathbf{B} \upharpoonright (\bigcup_{i \leq j} Nbd(3^{r-j}, b_i))$ under an isomorphism mapping a_i to b_i , for $1 \leq i \leq j$.

We first show that, for j = 1, the duplicator can ensure that the *j*-matching condition holds after the first round. Suppose that the spoiler selects a_1 from **A**. Let τ be the *d*-type of a_1 in **A**. Since **A** and **B** are (d, m)-equivalent, there is at least one point that has *d*-type τ in **B**. The duplicator selects one such point to be b_1 . Since $d = 3^{r-1}$ and j = 1, the *j*-matching condition is identical to the definition that a_1 and b_1 have the same *d*-type. (By symmetry, the same strategy works when we reverse the roles of **A** and **B**.)

We now show that if $1 \le j < r$ and if the *j*-matching condition holds, and the spoiler selects a_{j+1} from A, then the

duplicator can select b_{j+1} from **B** so that the (j+1)-matching condition holds (again, by symmetry, the same is true when we reverse the roles of **A** and **B**).

There are two cases.

Case 1. If

$$a_{j+1} \in \bigcup_{i \le j} Nbd(2 \cdot 3^{r-j-1}, a_i), \tag{3}$$

then

$$Nbd(3^{r-j-1}, a_{j+1}) \subseteq \bigcup_{i \leq j} Nbd(3^{r-j}, a_i)$$

Then the duplicator can select b_{j+1} to be the corresponding point of **B** (given by the isomorphism of the *j*-matching condition).

Case 2. If (3) fails, then let τ be the 3^{r-j-1} -type of a_{j+1} . Let l be the number of points in $R_A = \bigcup_{i \leq j} Nbd(2 \cdot 3^{r-j-1}, a_i)$ with 3^{r-j-1} -type τ . Let us denote by A' (resp. B') the structure on the left-hand side (resp. right-hand side) of the isomorphism in the *j*-matching condition. Now for every point $a \in R_A$, we have

$$Nbd(3^{r-j-1}, a) \subseteq \bigcup_{i \leq j} Nbd(3^{r-j}, a_i).$$

Therefore, for each point $a \in R_A$, the 3^{r-j-1} -type of a in A is the same as its 3^{r-j-1} -type in A'. Hence, l equals the number of points in R_A whose 3^{r-j-1} -type in A' is τ . Let $R_B = \bigcup_{i \leq j} Nbd(2 \cdot 3^{r-j-1}, b_i)$. By the *j*-matching condition, l also equals the number of points in R_B whose 3^{r-j-1} -type in B' is τ . Just as in the situation with A and A', it follows that lequals the number of points in R_B whose 3^{r-j-1} -type in B is τ .

We now show that l < m. Since $2 \cdot 3^{r-j-1} < 3^{r-1} = d$, it follows that $2 \cdot 3^{r-j-1} \le d-1$. Therefore, by Lemma 4.2, for each $i \le j$ the number of points in $Nbd(2 \cdot 3^{r-j-1}, a_i)$ is less than f^{d-1} . So j such neighborhoods have less than $j \cdot f^{d-1} < r \cdot f^{d-1} = m$ points altogether. Therefore, R_A has less than m points, so certainly l < m.

There are at least l+1 points in A with 3^{r-j-1} -type τ (namely, a_{j+1} , along with the *l* points in R_A with 3^{r-j-1} -type τ). Now A and B are (d, m)-equivalent, and hence, by Lemma 4.1, they are (d', m)-equivalent, where $d' = 3^{r-j-1} < d$. Therefore, since $l+1 \le m$, and since there are at least l+1 points in A with 3^{r-j-1} -type τ , there are also at least l+1 points in B with 3^{r-j-1} -type τ . In particular, there is some such point x outside of R_B , since R_B contains only l such points. Define b_{j+1} to be x. From the fact that (3) fails, it is easy to see that $A \upharpoonright \bigcup_{i \le j} Nbd(3^{r-j-1}, a_i)$ contains no point adjacent to a member of $A \upharpoonright Nbd(3^{r-j-1}, a_{j+1})$ (and similarly in B). It follows easily that the (j+1)matching condition holds. Since the duplicator's strategy guarantees that the *j*-matching condition holds for $1 \le j \le r$, in particular the *r*-matching condition holds. The *r*-matching condition says that $\mathbf{A} \upharpoonright \{a_1, ..., a_r\} \cong \mathbf{B} \upharpoonright \{b_1, ..., b_r\}$ under an isomorphism mapping a_i to b_i , for $1 \le i \le r$. So the duplicator wins. Therefore, $\mathbf{A} \sim_r \mathbf{B}$.

Since d in Theorem 4.3 depends only on r, and since two d-equivalent structures are (d, m)-equivalent for every m, the following corollary is immediate.

COROLLARY 4.4. Let r be a positive integer. There is a positive integer d such that whenever A and B are d-equivalent structures, then $A \sim_r B$.

We now describe Hanf's Lemma [Han65] (rewritten slightly to match our terminology). Hanf was not doing finite-model theory, so his lemma deals with both finite and infinite structures. Two structures **A** and **B** are *elementarily equivalent* if they agree on all first-order sentences (that is, $\mathbf{A} \models \sigma$ iff $\mathbf{B} \models \sigma$, for every first-order sentence σ).

LEMMA 4.5 (Hanf's Lemma). Assume that every neighborhood in A and B contains finitely many points. Then A and B are elementarily equivalent provided that, for each integer d and each d-type τ , either

1. both A and B have infinitely many points of d-type τ , or

2. A and **B** have the same finite number of points of d-type τ .

By the results in [Ehr61, Fra54], two structures A and B are elementarily equivalent precisely if $A \sim_r B$ for all r > 0. Thus, Theorem 4.3 is closely related to Hanf's Lemma, but neither seems to directly imply the other. In particular, Hanf's Lemma as stated is not useful in the context of finite structures, since two finite structures are elementarily equivalent iff they are isomorphic (discussion of this well-known fact can be found in [Fag93]). Our version of Hanf's Lemma is better suited for inexpressibility proofs for classes of finite structures, since it can show a bounded form of elementary equivalence, namely $A \sim_r B$, in cases where A and B are nonisomorphic finite structures.³

A result on graphs that is very similar to Theorem 4.3 appears in [Tho91, Lemma 4.1]. Theorem 4.3 is also related to a result by Gaifman [Gai82], who proved that in a precise sense, first-order logic talks only about neighborhoods. We can think of Hanf's Lemma, as well as

³ We now discuss how $\mathbf{A} \sim \mathbf{B}$ is a limited form of elementary equivalence. The *quantifier depth* $QD(\varphi)$ of a first-order sentence φ is defined recursively as follows: $QD(\varphi) = 0$ if φ is quantifier-free; $QD(\neg \varphi) = QD(\varphi)$; $QD(\varphi_1 \land \varphi_2) = \max\{QD(\varphi_1), QD(\varphi_2)\}; QD(\exists \varphi) = 1 + QD(\varphi)$. It turns out (and is closely related to results of Ehrenfeucht [Ehr61] and Fraissé [Fra54]) that $\mathbf{A} \sim \mathbf{B}$ iff \mathbf{A} and \mathbf{B} agree on all first-order sentences of quantifier depth r (that is, if σ is a first-order sentence of quantifier depth r, then $\mathbf{A} \models \sigma$ iff $\mathbf{B} \models \sigma$).

our versions, Theorem 4.3 and Corollary 4.4, as giving a library subroutine for the duplicator. Thus, instead of coming up with a winning strategy for a game, we can make use of the winning strategy given by the result.

As a warm-up for the proof in the next section, let us see how to use Corollary 4.4 to show that connectivity is not first-order. Assume that \mathscr{S} (the class of connected graphs) is first-order. Let r be given by Theorem 3.3(a). We obtain a contradiction by showing that the duplicator wins the r-game over \mathscr{S} . Find d as in Corollary 4.4. Let G_0 be a cycle with 4d nodes, and let G_1 be the disjoint union of two cycles, each with 2d nodes. It is easy to see that every point in G_0 and G_1 has the same d-type. Since G_0 and G_1 have the same number of points, and all with the same d-type, it follows that G_0 and G_1 are d-equivalent. By Corollary 4.4 and our choice of d, it follows that $G_0 \sim_r G_1$. Now the duplicator has a winning strategy in the r-game over \mathscr{S} : he selects $G_0 \in \mathscr{S}$ and $G_1 \in \overline{\mathscr{S}}$. Since, as we just showed, $G_0 \sim_r G_1$, the duplicator can now win.

5. APPLICATION TO CONNECTIVITY

Now we apply Theorem 3.4 and Corollary 4.4 to give a very simple proof that connectivity is not in monadic NP (and in particular, a proof that is much simpler than Fagin's original proof in [Fag75a]). Even though the details of our proof are not difficult, it is instructive to first outline the basic method, since a similar method might be applied to properties other than connectivity.

Let \mathcal{S} be the class of connected graphs. Assuming that connectivity is in monadic NP, we obtain a contradiction by giving, for all constants c and r, a winning strategy for the duplicator in the Ajtai-Fagin (c, r)-game over \mathcal{S} . The duplicator begins by choosing G_0 to be a sufficiently long cycle. After the spoiler colors G_0 with the c colors, the duplicator finds points α_p and α_q of G_0 such that α_p and α_q are sufficiently far apart and such that, intuitively, the coloring of points near to α_p looks the same as the coloring of points near to α_q . Here, the precise definition of "sufficiently far apart" and "near" both depend on the parameter d given by Corollary 4.4. The duplicator then "pinches" G_0 together at the points α_p and α_q to split G_0 into two disjoint cycles. This pair of disjoint cycles is the duplicator's choice for G_1 , and the coloring of G_1 is inherited from G_0 . Thus, the duplicator's coloring strategy is trivial. It then follows that, for every d-type τ , the graphs G_0 and G_1 have exactly the same number of points with d-type τ . Corollary 4.4 can then be applied to conclude that $G_0 \sim_r G_1$. In the proof of the following, the details are filled in.

THEOREM 5.1. Connectivity is not in monadic NP.

Proof. Let \mathscr{G} be the class of connected graphs and assume that \mathscr{G} is in monadic NP. Let c, r be given by

Theorem 3.4. We obtain a contradiction by showing that the duplicator wins the Ajtai-Fagin (c, r)-game over \mathcal{S} .

Let *d* be given by Corollary 4.4 for this *r*. The duplicator chooses G_0 to be a directed cycle of length *n*, for a sufficiently large *n*. Let $\alpha_0, \alpha_1, ..., \alpha_{n-1}$ denote the points in order around the cycle, so that there is an edge from α_i to α_{i+1} for $0 \le i < n$. Here and subsequently, subscripts are reduced modulo *n* to belong to the interval [0, n-1].

The spoiler now colors G_0 with c colors. Let $\chi(\alpha_i)$ denote the color of α_i . Assuming that $n \ge 2d$, the d-type of the point α_i in the resulting structure is fully described by the following vector of 2d - 1 colors:

$$\langle \chi(\alpha_{i-(d-1)}), ..., \chi(\alpha_{i-1}), \chi(\alpha_i), \chi(\alpha_{i+1}), ..., \chi(\alpha_{i+(d-1)}) \rangle.$$

The number of possible *d*-types is some constant, depending on *c* and *d*, but not on *n*. So it is clear that, for *n* sufficiently large, there must be at least 4*d* points with the same *d*-type. Therefore, there must exist points α_p and α_q that have the same *d*-type and are at least distance 2*d* apart (that is, $\alpha_p \notin Nbd(2d, \alpha_q)$).

The duplicator now forms G_1 , a pair of disjoint directed cycles, by pinching G_0 together at the points α_p and α_q (see Fig. 1). More precisely, let G_1 be a structure with universe consisting of *n* distinct points β_0 , β_1 , ..., β_{n-1} . There is an edge from β_i to β_{i+1} for all *i* with $0 \le i < n$, $i \ne p$, and $i \ne q$,



there is an edge from β_p to β_{q+1} , and there is an edge from β_q to β_{p+1} . There are no other edges. The duplicator's coloring of G_1 is given by $\chi(\beta_i) = \chi(\alpha_i)$ for all *i*.

Note that each component cycle of G_0 or G_1 contains at least 2*d* points, since α_p and α_q are at least distance 2*d* apart. Since also α_p and α_q have the same *d*-type, it follows that α_i and β_i have the same *d*-type for all *i*, so G_0 and G_1 are *d*-equivalent. It follows from Corollary 4.4 that $G_0 \sim_r G_1$, so the duplicator wins.

It is instructive to see why the use of the Ajtai-Fagin (c, r)-game, as opposed to the original (c, r)-game, is important for our proof. The choice of G_1 depends on the coloring by the spoiler of G_0 . Our proof would not work if, as in the original (c, r)-game, the duplicator were required to select G_0 and G_1 before the spoiler colors G_0 .

We note that Hájek [Háj75], independently of Fagin [Fag75a] but somewhat later, also proved that connectivity is not in monadic NP. Hájek's proof uses nonstandard analysis and semisets. Interestingly enough, Hájek's proof involves splitting a cycle, just as our proof does.

6. BUILT-IN RELATIONS OF MODERATE DEGREE

In this section the result that connectivity is not in monadic NP is extended to the case where sentences are allowed to contain built-in relations of moderate (i.e., sufficiently small) degree. This gives, in particular, the result of de Rougemont [dR87] that connectivity is not in monadic NP in the presence of a built-in successor relation. Let $V_n = \{v_0, v_1, ..., v_{n-1}\}$ be a universe of size *n*. A particular collection of built-in relations is specified by a language $\{P_1, ..., P_s\}$ and, for each $n \ge 1$ and $1 \le i \le s$, an interpretation $\hat{P}_{n,i}$ of P_i as a relation on V_n . Let \mathbf{P}_n denote the structure with domain V_n and relations $\hat{P}_{n,1}, ..., \hat{P}_{n,s}$. Let f(n) be the maximum degree of a point v in the structure \mathbf{P}_n . (Recall from Section 4 that the degree of v is the number of points adjacent to v but not equal to v, where v and v' are adjacent (in \mathbf{P}_n) if they both belong to the same tuple in $\hat{P}_{n,i}$ for some i.) We say that the built-in relations have moderate degree if $f(n) = (\log n)^{o(1)}$, that is, if there is a function $\sigma(n)$ with $\lim_{n\to\infty} \sigma(n) = 0$ such that $f(n) \leq (\log n)^{\sigma(n)}$ for all n. (Although the base of the logarithm is not important, say that the base is 2 for definiteness.)

The next theorem is our main result.

THEOREM 6.1. Connectivity is not in monadic NP, even in the presence of built-in relations of moderate degree.

Proof. The proof is similar to the proof of Theorem 5.1, although the details are more complicated. Fix some collection of built-in relations as above. Assume that the class \mathscr{S} of connected graphs is in monadic NP using these built-in relations. Let c, r be given by Theorem 3.4. Although Theorem 3.4 is stated for the case where there are no built-in

relations, it is clear how the Ajtai-Fagin game generalizes to the case of built-in relations, and Theorem 3.4 remains true in this case. For example, in the first step, the duplicator extends the "built-in" structure P_n by choosing an interpretation E_0 for the edge relation of a connected graph on V_n . As before, the duplicator will choose E_0 to be a directed cycle on V_n , henceforth called a *full cycle*. As before, we want to show that the duplicator can choose the full cycle E_0 such that, no matter how the spoiler colors the points, the duplicator can split the cycle into two disjoint cycles in such a way that Corollary 4.4 applies to the two structures. The new difficulty is that the built-in relations impose an additional structure on the points, and we know nothing about this structure, other than that it has moderate degree. (It should perhaps be noted that the duplicator cannot choose E_0 arbitrarily. As a simple example, suppose that one of the built-in relations is itself a full cycle. If the duplicator chooses E_0 to coincide with this built-in relation then the spoiler will always win no matter how the duplicator splits the cycle.) We show that if the duplicator chooses a cycle at random, then with high probability the chosen cycle will work to defeat the spoiler. In particular, since the probability is positive, this shows that there exists a choice that the duplicator can make that is guaranteed to defeat the spoiler.

We first outline the duplicator's winning strategy in the Ajtai-Fagin (c, r)-game. Let d be given by Corollary 4.4 for this r. The duplicator chooses a sufficiently large n. Abbreviate $V = V_n$ and $\mathbf{P} = \mathbf{P}_n$. If E is a binary (edge) relation, let $\mathbf{P}_n(E)$ (abbreviated $\mathbf{P}(E)$) denote the structure on the universe V_n with relations $\hat{P}_{n,1}, ..., \hat{P}_{n,s}, E$. A key concept in the proof is the notion of a point being "good" for a full cycle E. Let $\alpha_0, \alpha_1, ..., \alpha_{n-1}$ denote the points in order around E. Informally, we say that α_i is good for E if a sufficiently large neighborhood around α_i can be partitioned into a left part L and a right part R such that $\alpha_i \in L, \alpha_{i+1} \in R$, the only adjacency between a point of L and a point of R is the adjacency between α_i and α_{i+1} , and, moreover, this adjacency occurs only in the cycle E, not in any of the built-in relations. Intuitively, the left part and the right part are "unrelated," except for the single cycle edge from α_i to α_{i+1} . We show that if the duplicator chooses a full cycle E uniformly at random, then he can expect that "many" points will be good for E (this is one place where we use the assumption of moderate degree). In particular, there must exist a full cycle E_0 for which "many" points are good. The duplicator begins by choosing the structure $P(E_0)$. After the spoiler colors the points of $P(E_0)$ with the c colors, let A denote the resulting colored structure. Since there are a large number of good points, and since the number of possible 2*d*-types can be shown to be much smaller than n(this is another place where we use the assumption of moderate degree), we can find points α_p and α_q such that α_p and α_q are good, α_p and α_q have the same 2*d*-type in **A**, and

 α_p and α_q are sufficiently "far apart" in A. Similar to the proof of Theorem 5.1, the duplicator then forms E_1 , consisting of two disjoint directed cycles, by removing two edges from E_0 and adding two other edges to E_0 , as shown in Figure 1. Let B be the resulting structure. That is, B is P, with all points colored the same as in A, but extended by the edge relation E_1 instead of E_0 . It will then follow from the properties of α_p and α_q that, for every *d*-type τ , the structures A and B have exactly the same number of points with d-type τ . In particular, since α_p and α_q are good and are sufficiently far apart, the left and right parts associated with α_p and the left and right parts associated with α_q are pairwise unrelated in A, except for the cycle edge from α_p to α_{p+1} and the cycle edge from α_q to α_{q+1} . Since, in addition, α_p and α_q have the same 2d-type, we can argue that the d-type of every point is not changed when these two cycle edges are replaced by edges from α_p to α_{q+1} and from α_q to α_{p+1} . By Corollary 4.4, $\mathbf{A} \sim_r \mathbf{B}$. So the duplicator wins the Ajtai-Fagin (c, r)game over \mathscr{S} .

In what follows, we will be dealing with neighborhoods in a variety of different structures. For a structure S, let Nbd(S; d, v) denote the neighborhood Nbd(d, v) in the structure S.

We begin to fill in the details by giving the formal definition of a point being good for a full cycle *E*. Let $\alpha_0, \alpha_1, ..., \alpha_{n-1}$ denote the points in order around *E*; that is, $\langle \alpha_i, \alpha_j \rangle \in E$ iff $j \equiv i + 1 \mod n$. Let E(i) be the relation *E* with the tuple $\langle \alpha_i, \alpha_{i+1} \rangle$ removed. (As before, subscripts are reduced modulo *n*.) The point α_i is good for *E* if *Nbd*(**P**(*E*); 2*d*, α_i) can be partitioned into a *left part L* and a *right part R* (that is, $L \cup R = Nbd(\mathbf{P}(E); 2d, \alpha_i)$ and $L \cap R = \emptyset$), such that $\alpha_i \in L, \alpha_{i+1} \in R$, and, for every $u \in L$ and $u' \in R$, the points *u* and *u'* are not adjacent in **P**(*E*(*i*)). Any pair $\langle L, R \rangle$ having these properties is said to be a good partition for α_i in *E*.

The next step is to show that there is a full cycle E_0 for which at least half of the points are good. This is done by a probabilistic argument. By a randomly chosen full cycle we mean one chosen from the uniform distribution where every full cycle has the same probability (1/(n-1)!) of being chosen. The key step is to prove the following:

CLAIM 1. Assume that n is sufficiently large. Fix a point v. If E is a randomly chosen full cycle, then the probability that v is good for E is at least $\frac{1}{2}$.

From this it follows immediately that, for a randomly chosen full cycle E, the expected number of good points is at least n/2. So there must be some E_0 with at least n/2 good points. The following terminology will be useful in proving Claim 1. If $F \subseteq V \times V$, say that F is legal if $F \subseteq F'$ for some full cycle F'. Thus, F is legal if it consists of a subset of the edges of a full cycle. If $\langle v, v' \rangle \in F$ where F is legal, v is the left neighbor of v' in F, and v' is the right neighbor of v in F. The point v is *full in* F if v has a left neighbor and a right neighbor in F. The point v is *deficient in* F if it is not full.

To prove Claim 1, fix an arbitrary $v \in V$ and consider the random procedure RandomCycle(v). The output is a pair $\langle E, b \rangle$ where E is a full cycle and $b \in \{0, 1\}$. The output b=0 means that the procedure was successful in producing an E for which v is good. The procedure RandomCycle(v)is described precisely below. In outline, the procedure operates as follows. The procedure has n steps and constructs a full cycle by randomly adding one edge at each step. The procedure has two phases. In the first phase (labelled 2(a) below), the procedure (repeatedly and deterministically) chooses a deficient point v_i in the neighborhood of radius 2d about v, and randomly adds a cycle edge incident on v_i by randomly choosing the other endpoint z of the edge from among the z's that preserve legality of the set of chosen edges; this continues until no point in this neighborhood is deficient. (This neighborhood grows as edges are added.) During this phase, v remains good for the cycle being chosen provided that the randomly chosen endpoint of each added edge does not belong to a set of "bad" points that lie "too close" to v. Intuitively, this works for the following reason. Let v' be the point reached by following the cycle edge from v. If v is not good for the chosen cycle, then, for every way of partitioning the neighborhood of radius 2d about v into a left part L containing v and a right part R containing v', there will be an adjacency between L and R other than the adjacency caused by the cycle edge $\langle v, v' \rangle$; this implies the existence of a "short" path of adjacencies from v to v' not involving the cycle edge $\langle v, v' \rangle$. If the randomly chosen endpoints lie sufficiently far from v, then such short paths of adjacencies are avoided. A detailed proof of this is given in the proof of Claim 1.2 below. Moreover, the set of bad points is sufficiently small, and the first phase lasts for sufficiently few steps, that all points chosen in the first phase are not bad with high probability. A detailed proof of this is given in the proof of Claim 1.3 below. During the second phase (labelled 2(b) below), any remaining deficient points are made full by randomly choosing cycle edges incident on these points. The choices made during the second phase do not affect v being good for the chosen cycle, since these edges are all incident on points outside the neighborhood of radius 2d about v. The procedure uses S_t to denote the set of edges chosen during the first t-1 steps, and it uses B, to denote the set of points that are "bad" choices for the randomly chosen endpoint of the edge added at step t. In general, B_t is the neighborhood of radius 4d about v in the structure **P** extended by the set of cycle edges chosen so far.

RANDOMCYCLE(v):

- 1. Set $S_1 = \emptyset$, $B_1 = Nbd(\mathbf{P}; 4d, v)$, and b = 0.
- 2. Do the following for t = 1, 2, ..., n:

(a) If there is a point $v_i \in Nbd(\mathbf{P}(S_t); 2d, v)$ that is deficient in S_t , let v_i be such a point with smallest i. (We choose the minimum i just to be definite; this choice is not important.) If v_i does not have a right neighbor, choose a right neighbor z for v_i at random; that is choose a point z uniformly at random from those points z such that $S_t \cup \{\langle v_i, z \rangle\}$ is legal. Set $S_{t+1} = S_t \cup \{\langle v_i, z \rangle\}$. If v_i does have a right neighbor, then randomly choose a left neighbor z for v_i in a similar way. In either case, if $z \in B_t$ then set b = 1. Set $B_{t+1} =$ $Nbd(\mathbf{P}(S_{t+1}); 4d, v)$.

(b) If there is no point $v_i \in Nbd(\mathbf{P}(S_i); 2d, v)$ that is deficient in S_i , then let v_i be a point with smallest i that is deficient in S_i . Choose a random right or left neighbor z for v_i as above, and let S_{i+1} be S_i together with the new edge $\langle v_i, z \rangle$ or $\langle z, v_i \rangle$. In this case, b does not change, and we can (arbitrarily) take $B_{i+1} = B_i$.

3. Output $\langle S_{n+1}, b \rangle$.

To see why the choice of B_{t+1} is not important in case 2(b), note that if at some step t we have that no point of $Nbd(\mathbf{P}(S_t); 2d, v)$ is deficient, then $Nbd(\mathbf{P}(S_j); 2d, v) = Nbd(\mathbf{P}(S_t); 2d, v)$ for all j > t. Thus, case 2(a) will occur at all steps up to some step t_0 , and case 2(b) will occur at all steps thereafter (where t_0 may depend on the random choices made by the procedure). This fact is used again below.

Claim 1 follows immediately from the following three claims about the output $\langle E, b \rangle$ of RandomCycle(v), for each fixed v.

CLAIM 1.1. The distribution of E is the uniform distribution on full cycles.

CLAIM 1.2. If b = 0, then v is good for E.

CLAIM 1.3. $\Pr[b=0] \ge 1/2$, for all sufficiently large n.

Proof of Claim 1.1. First observe that, at each step t with $1 \le t < n$, if v_i does not have a right (resp., left) neighbor, there are precisely n-t points z that are candidates for a right (resp., left) neighbor of v_i . To see this, note that if t = 1 there are n - 1 candidates (namely, every point except v_i itself), and every time we add a new edge to E the number of candidates decreases by one. So there are a total of (n-1)! different choices that the procedure can make at all *n* steps. Next observe that, for every full cycle \hat{E} , there is a way for the procedure to make choices so that the output will be $\langle \hat{E}, b \rangle$ for some b. That is, when the procedure is choosing a right (left) neighbor for v_i in step 2, it chooses the right (left) neighbor of v_i in \hat{E} . It follows that there is a one-to-one correspondence between the (n-1)!full cycles and the (n-1)! choices, so each full cycle is produced with the same probability (1/(n-1)!). This completes the proof of Claim 1.1.

Proof of Claim 1.2. Suppose $\langle E, 0 \rangle$ is produced. As above, let $\alpha_0, ..., \alpha_{n-1}$ be the points in order around *E*. Let

i be such that $v = \alpha_i$. Recall that E(i) is *E* with the tuple $\langle \alpha_i, \alpha_{i+1} \rangle$ removed. It is useful to view the adjacency relationship in $\mathbf{P}(E)$ as an undirected graph *H*. The vertex set of *H* is *V*, and there is an undirected edge (u, u') iff $u \neq u'$ and *u* and *u'* are adjacent in $\mathbf{P}(E)$. Define the adjacency graph H(i) similarly for the structure $\mathbf{P}(E(i))$. Observe that α_i and α_{i+1} are not adjacent in the built-in structure \mathbf{P} . For if they were, we would have both α_i and α_{i+1} in every bad set B_t for all $t \ge 1$, so adding the edge $\langle \alpha_i, \alpha_{i+1} \rangle$ to *E* at some step *t* would cause *b* to be set to 1. So the only difference between *H* and H(i) is that (α_i, α_{i+1}) is not an edge of H(i), whereas it is an edge of *H*.

Let $L = Nbd(\mathbf{P}(E(i)); 2d, \alpha_i)$ and $R = Nbd(\mathbf{P}(E(i)); 2d-1, \alpha_{i+1})$. Note that $L \cup R = Nbd(\mathbf{P}(E); 2d, \alpha_i)$, since we reach a point u of $Nbd(\mathbf{P}(E); 2d, \alpha_i)$ either by following a path in H of length at most 2d - 1 from α_i to u not using the edge (α_i, α_{i+1}) , or by first following the edge (α_i, α_{i+1}) and then following a path of length at most 2d - 2 from α_{i+1} to u. Obviously, $\alpha_i \in L$ and $\alpha_{i+1} \in R$.

To complete the proof that $\langle L, R \rangle$ is a good partition for $v = \alpha_i$ in E (and, therefore, that v is good for E), we must show that, for every $u \in L$ and $u' \in R$, the points u and u' are not adjacent in $\mathbf{P}(E(i))$. (Since a point is adjacent to itself, this implies, in particular, that $L \cap R = \emptyset$.) Assume for contradiction that $u \in L$, $u' \in R$, and u and u' are adjacent in $\mathbf{P}(E(i))$. Recall that (α_i, α_{i+1}) is not an edge of H(i). We now show that there is a path from α_i to α_{i+1} in H(i), having length at least 2 and at most 4d-2. (The length is at least 2, since (α_i, α_{i+1}) is not an edge of H(i).) Beginning at α_i , the path first follows a path of length at most 2d-1 to u, then follows a path of length at most one to u', and then follows a path of length at most 2d-2 to α_{i+1} . So there is also a simple (i.e., having no repeated vertices) path from α_i to α_{i+1} in H(i) of length at least 2 and at most 4d-2. Since the edge (α_i, α_{i+1}) appears in *H*, there is a simple cycle *C* in H, of length at least 3 and at most 4d-1, passing through α_i and α_{i+1} . Note that there is a pair $\langle y, y' \rangle$, namely, $\langle \alpha_i, \alpha_{i+1} \rangle$, such that y and y' are connected by an edge of C and $\langle y, y' \rangle \in E$. Among all the pairs $\langle y, y' \rangle$ meeting these two conditions, consider the one that Random-Cycle(v) added last to E. Say that this happens at step t. Since all the adjacencies in C, except possibly (y, y'), occur in the structure $\mathbf{P}(S_t)$, and since $v = \alpha_t$ is on the cycle C, we have y, $y' \in Nbd(\mathbf{P}(S_t); 4d, v) \subseteq B_t$, and at least one of y or y' belongs to $Nbd(\mathbf{P}(S_t); 2d, v)$. Since both y and y' are deficient in S, (since the edge $\langle y, y' \rangle$ is added at step t), it must be that case 2(a) occurs at step t. Since y, $y' \in B_t$, adding the edge $\langle y, y' \rangle$ to E at step t would cause b to be set to 1. This contradiction completes the proof of Claim 1.2.

Proof of Claim 1.3. Recall that f(n) is the maximum degree of a point in the built-in structure \mathbf{P}_n , and that $f(n) = (\log n)^{o(1)}$. Hence, (f(n) + 2) is the maximum degree of a point in $\mathbf{P}_n(E)$. Let $\mu(n) = (f(n) + 2)^{2d}$. By Lémma 4.2,

it follows that $\mu(n)$ is an upper bound on the number of points in a neighborhood of radius 2d in a structure where every point has degree at most f(n) + 2. So Nbd($\mathbf{P}(S_t)$; 2d, v) contains at most $\mu(n)$ points for any t. Since $t_1 < t_2$ implies $Nbd(\mathbf{P}(S_t); 2d, v) \subseteq Nbd(\mathbf{P}(S_t); 2d, v)$, after the procedure has added $2\mu(n)$ edges to E (i.e., at step $t = 2\mu(n) + 1$, no point of Nbd($\mathbf{P}(S_t)$; 2d, v) will be deficient in S_t . Therefore, letting t_0 be such that case 2(a) occurs at steps 1, 2, ..., t_0 , and case 2(b) occurs thereafter, we have $t_0 \leq 2\mu(n)$. Since $f(n) = (\log n)^{o(1)}$, we also have $\mu(n) =$ $(\log n)^{o(1)}$. Similarly, B_t contains at most $(f(n)+2)^{4d} =$ $(\mu(n))^2$ points for every t. At each step $t \leq 2\mu(n)$, there are at least $n - 4\mu(n)$ "legal" choices for z in 2(a), since every point that is not an endpoint of an edge in S_i is a legal choice. So at each step $t \leq 2\mu(n)$, the probability that z is chosen in B_r , and therefore the probability that b is set to 1, is at most $(\mu(n))^2/(n-4\mu(n))$. Therefore, the probability that b is set to 1 during the first $2\mu(n)$ steps is at most

$$2\mu(n)\left(\frac{(\mu(n))^2}{n-4\mu(n)}\right).$$
 (4)

Since $\mu(n) = (\log n)^{o(1)}$, the above expression gets arbitrarily small as *n* gets arbitrarily large. Thus, the probability that the final value of *b* is 0 is at least 1/2 for sufficiently large *n*. This completes the proof of Claim 1.3.

Having proved Claim 1, we can now complete the proof of Theorem 6.1, following the outline above. The duplicator first chooses the structure $P(E_0)$ where E_0 is a full cycle for which at least n/2 points are good. The spoiler then colors the points of $P(E_0)$ with the c colors. Let A denote the resulting structure.

We need the following simple upper bound on the number of 2d-types in A.

CLAIM 2. For all sufficiently large n, there are at most \sqrt{n} different 2d-types in A.

Proof. As above, $\mu(n) = (f(n) + 2)^{2d}$ is an upper bound on the number of points in any neighborhood of radius 2d in A. For appropriate constants a and b, an upper bound on the number of 2d-types in A is $N(n) = 2^{b(\mu(n))^a}$. To see this, let a be the maximum arity of the built-in relations and the edge relation. In any set S containing at most $\mu(n)$ points, there are at most $(\mu(n))^a$ distinct a'-tuples if $a' \leq a$, so there are at most $2^{(\mu(n))^a}$ interpretations of an a'-ary predicate on S. So for a collection of b such predicates, there are at most N(n) different interpretations. Since $\mu(n) = (\log n)^{o(1)}$, we have $N(n) \leq \sqrt{n}$ for sufficiently large n. This proves Claim 2.

The next claim states that we can find points α_p and α_q having certain properties. These properties then allow us to prove that, if **B** is obtained from **A** by splitting the cycle at α_p and α_q , then **A** and **B** are *d*-equivalent.

CLAIM 3. If n is sufficiently large, then there are points α_p and α_q such that (1) α_p and α_q are good for E_0 , (2) α_p and α_q have the same 2d-type in A, and (3) $\alpha_q \notin Nbd(A; 4d, \alpha_p)$.

Proof. Since there are at least n/2 points good for E_0 , and since there are at most \sqrt{n} different 2*d*-types in **A** (by Claim 2 for *n* sufficiently large), there must be a set of at least $\sqrt{n/2}$ points good for E_0 that have the same 2*d*-type τ in **A**. Let α_p be an arbitrary point with 2*d*-type τ that is good for E_0 . Since the size of $Nbd(\mathbf{A}; 4d, \alpha_p)$ is $(\log n)^{o(1)}$, there must be a point α_q with 2*d*-type τ that is good for E_0 and that is not in $Nbd(\mathbf{A}; 4d, \alpha_p)$. This proves Claim 3.

For $i \in \{p, q\}$, let $\langle L_i, R_i \rangle$ be a good partition for α_i in E_0 . The relevant part of **A** is shown in the top half of Fig. 2. By choice of α_p and α_q , if *u* is a point in one of the sets L_p, R_p, L_q, R_q , and *u'* is a point in a different one of these sets, then *u* and *u'* are not adjacent in the built-in structure **P**. Moreover, such a *u* and *u'* are adjacent in E_0 iff $\{u, u'\} = \{\alpha_p, \alpha_{p+1}\}$ or $\{u, u'\} = \{\alpha_q, \alpha_{q+1}\}$.

The duplicator now forms a structure **B** as follows. With regard to the built-in relations, let **P'** be an isomorphic copy of **P** on the universe $\{\beta_0, \beta_1, ..., \beta_{n-1}\}$ under the isomorphism mapping α_i to β_i for each *i*. Define the edge relation E_1 by $\langle \beta_i, \beta_{i+1} \rangle \in E_1$ for all *i* with $i \neq p$ and $i \neq q$, and $\langle \beta_p, \beta_{q+1} \rangle \in E_1$, and $\langle \beta_q, \beta_{p+1} \rangle \in E_1$. Let $\mathbf{B} = \mathbf{P}'(E_1)$. (See the bottom half of Fig. 2.)

To complete the proof we argue that, for every j, the points α_j and β_j have the same *d*-type. This implies that **A** and **B** are *d*-equivalent, so, by Corollary 4.4, the duplicator wins. In the rest of the proof, neighborhoods of points α_i are



taken in A, and neighborhoods of points β_i are taken in B. Define the "identity" map *i* by $\iota(\alpha_i) = \beta_i$ for all *i*. Let L'_p (resp., R'_p) be the image under *i* of L_p (resp., R_p). Define L'_q and R'_q similarly. (Again see Fig. 2.) Fix an arbitrary *j* with $0 \le j < n$. If $\alpha_j \notin Nbd(d, \alpha_p)$ and $\alpha_j \notin Nbd(d, \alpha_q)$, then clearly α_j and β_j have the same *d*-type since these *d*-types do not involve any of the edges that were changed in going from E_0 to E_1 . (So a suitable restriction of *i* gives an isomorphism between $Nbd(d, \alpha_j)$ and $Nbd(d, \beta_j)$.) Say that $\alpha_j \in Nbd(d, \alpha_p)$. (The case $\alpha_j \in Nbd(d, \alpha_q)$ is completely symmetric and is omitted.) Note that

$$Nbd(d, \alpha_i) \subseteq Nbd(2d, \alpha_n) = L_n \cup R_n.$$

Say that $\alpha_j \in L_p$. (The case $\alpha_j \in R_p$ is similar and is omitted.) If $\alpha_{p+1} \notin Nbd(d, \alpha_j)$, then $Nbd(d, \alpha_j)$ is contained completely in L_p , so it is again obvious that α_j and β_j have the same *d*-type (with *i* giving the isomorphism). So assume that $Nbd(d, \alpha_j)$ contains both α_p and α_{p+1} .

First, the map ι restricted to $Nbd(d, \alpha_j) \cap L_p$ gives an isomorphism

$$\mathbf{A} \upharpoonright (Nbd(d, \alpha_i) \cap L_n) \cong \mathbf{B} \upharpoonright (Nbd(d, \beta_i) \cap L'_n)$$

mapping α_j to β_j and mapping α_p to β_p . Call this isomorphism π_L . Since $\mathbf{A} \upharpoonright Nbd(2d, \alpha_p) \cong \mathbf{A} \upharpoonright Nbd(2d, \alpha_q)$ under an isomorphism mapping α_p to α_q , it follows that $\mathbf{A} \upharpoonright R_p \cong \mathbf{A} \upharpoonright R_q$ under an isomorphism π_1 mapping α_{p+1} to α_{q+1} . (It also follows that $\mathbf{A} \upharpoonright L_p \cong \mathbf{A} \upharpoonright L_q$ under an isomorphism mapping α_p to α_q , although we do not need this isomorphism for the present case.) The restriction of *i* to R_q gives $\mathbf{A} \upharpoonright R_q \cong \mathbf{B} \upharpoonright R'_q$ under an isomorphism π_2 mapping α_{q+1} to β_{q+1} . So the composition of π_1 and π_2 is an isomorphism $\mathbf{A} \upharpoonright R_p \cong \mathbf{B} \upharpoonright R'_q$ mapping α_{p+1} to β_{q+1} . Restricting this latter isomorphism to $Nbd(d, \alpha_j) \cap R_p$ gives π_R , an isomorphism

$$\mathbf{A} \upharpoonright (Nbd(d, \alpha_j) \cap R_p) \cong \mathbf{B} \upharpoonright (Nbd(d, \beta_j) \cap R'_q)$$

mapping α_{p+1} to β_{q+1} . Since the only adjacency between a point of $Nbd(d, \alpha_j) \cap L_p$ and a point of $Nbd(d, \alpha_j) \cap R_p$ is the adjacency between α_p and α_{p+1} in E_0 , and since the only adjacency between a point of $Nbd(d, \beta_j) \cap L'_p$ and a point of $Nbd(d, \beta_j) \cap R'_q$ is the adjacency between β_p and β_{q+1} in E_1 , it follows that the union of π_L and π_R is an isomorphism $\mathbf{A} \upharpoonright Nbd(d, \alpha_j) \cong \mathbf{B} \upharpoonright Nbd(d, \beta_j)$ mapping α_j to β_j . So α_j and β_j have the same *d*-type.

As an immediate corollary of Theorem 6.1, we obtain the following result, originally proved by de Rougemont [dR87].

COROLLARY 6.2. Connectivity is not in monadic NP in the presence of a built-in successor relation.

Since it is an open question⁴ whether Theorem 6.1 remains true with a degree bound larger than $(\log n)^{o(1)}$, it is instructive to see where our proof of Theorem 6.1 breaks down when the degree bound is increased. The first place where we used the degree bound was in the proof of Claim 1.3. This proof remains valid with the larger degree bound $n^{o(1)}$, since the probability in (4) still gets arbitrarily small as *n* increases. (To see this, note that $\mu(n) = n^{o(1)}$ if the degree is at most $f(n) = n^{o(1)}$.) The second place where we used the degree bound (the more critical place) was in the proof of Claim 2 placing an upper bound on the number of 2d-types in A. If a built-in relation has degree $f(n) = (\log n)^{\varepsilon}$ for some constant $\varepsilon > 0$, then, since our proof must work for an arbitrary d, the size of neighborhoods of radius 2d can exceed $\log n$ for large enough d. Even for a single 2-ary relation, the number of non-isomorphic structures on $\log n$ points exceeds n (for all sufficiently large n), so we cannot argue as before that there must be two points with the same 2d-type to use in splitting the cycle. (Although it was convenient to use 2d-types rather than *d*-types in the proof, this difficulty holds for *d*-types as well.)

Regarding the problem of showing that connectivity is not in monadic NP in the presence of arbitrary built-in relations, it should be noted that monadic NP with arbitrary built-in relations has a natural circuit characterization in terms of *nondeterministic* AC^{0} where the number of bits of nondeterminism is linear in the size of the universe. Intuitively, kn bits of nondeterminism correspond to kexistentially-quantified monadic relations. This characterization follows easily from the equivalence of AC⁰ and first-order logic with arbitrary built-in relations [Imm87]. More precisely, a class & of graphs that is closed under isomorphism is in monadic NP in the presence of some collection of built-in relations iff there is a positive integer kand a family $H = \{h_n \mid n \ge 1\}$ of functions such that, for each n, (i) h_n is a Boolean-valued function of n^2 Boolean "edge inputs" $\{x_{i,j} \mid 0 \le i, j \le n-1\}$ and kn Boolean "nondeterministic inputs" $\{y_i | 1 \le i \le kn\}$, (ii) the family H belongs to (nonuniform) AC^0 (i.e., there is a polynomial p and a constant d such that, for each n, the function h_n is computed by some circuit of size p(n) and depth d containing and-gates, or-gates, and not-gates), and (iii) for every graph G on n points $\{0, 1, ..., n-1\}$, if the edge inputs are set according to the edges of G (i.e., $e_{i,j} = 1$ iff there is an edge from point i to point j in G), then G is in the class \mathscr{C} iff

$$(\exists y_1 \cdots \exists y_{kn})(h_n(e_{0,0}, ..., e_{n-1,n-1}, y_1, ..., y_{kn}) = 1).$$

We note that Ajtai [Ajt83] has shown that the property "the number of edges of G is even" is not in monadic NP in the presence of arbitrary built-in relations.

⁴ Note added in proof. Schwentick [Sch94a] has recently extended this result to hold with a degree bound of $n^{o(1)}$

7. CONCLUSIONS AND OPEN QUESTIONS

We have given a new, simple proof of Fagin's result that connectivity is not in monadic NP. Furthermore, we have extended this result, and extended de Rougemont's result that connectivity is not in monadic NP even in the presence of successor, by showing that connectivity is not in monadic NP even in the presence of built-in relations of moderate degree, that is, degree $(\log n)^{o(1)}$.

Our proofs combined three techniques. First, we used a technique based on work of Hanf which permits the Ehrenfeucht-Fraïssé game part of the proof to be replaced by a combinatorial argument, counting d-types. Second, we made use of the Ajtai-Fagin (c, r)-game in place of the original (c, r)-game. Finally, we made use of another technique, also introduced by Ajtai and Fagin, of playing Ehrenfeucht-Fraïssé games over random structures. Note that the first two techniques were sufficient to enable us to obtain an almost trivial proof that connectivity is not in monadic NP. It is likely that the above methods can be applied to show that other graph properties are not in monadic NP. For example, we can show that non-3colorability is not in monadic NP. (Recall from Section 2 that 3-colorability is in monadic NP.) The proof follows the same outline as the proof for connectivity, although the graph G_0 is more complicated. Cosmadakis [Cos93] independently proved that non-3-colorability is not in monadic NP; moreover, he shows that this holds in the presence of a built-in successor relation. His proof works by giving a firstorder reduction from connectivity to non-3-colorability that expands the size of the universe at most linearly. Cosmadakis [Cos93] also uses this reduction method to show that several other problems in monadic co-NP are not in monadic NP, even in the presence of a built-in successor relation.

While an exact characterization of the monadic NP graph properties may be too much to hope for, some type of "general" result, showing that a large number of graph properties are not in monadic NP, seems feasible at this point. In analogy, there are results that establish NP-completeness for large classes of graph properties, for example, results of [LY80].

An open problem is to extend Theorem 6.1 to built-in relations of larger degree.⁵ Another interesting open problem in the area of Σ_1^1 inexpressibility results for graph properties is to extend beyond the monadic case. As noted in [Fag75b, Fag93], even the following is open: Is there a property of graphs that can be expressed by some (nonmonadic) Σ_1^1 sentence (equivalently, a property that belongs to NP), but that cannot be expressed by a sentence of the form $\exists Q\psi$ where Q is a single binary relation and ψ is first-order? Extending beyond the monadic case is an

⁵ But see Footnote 4.

important direction if the connection between "logic and complexity" is to have an impact on questions in computational complexity such as the NP = co-NP question. We believe that developing our descriptive complexity toolkit is a useful and necessary step, if we are ever to make progress on such difficult questions through descriptive complexity techniques. Evidence that we are moving in the right direction is that our tools are now powerful enough (1) to give a very simple proof of a result that used to have only a hard proof, and (2) to enable us to prove a new result, where we allow a large class of built-in relations.

ACKNOWLEDGMENT

We are grateful to Phokion Kolaitis for useful comments on a previous draft of this paper.

Received February 24, 1992; final manuscript received May 24, 1994

REFERENCES

- [AF90] Ajtai, M., and Fagin, R. (1990), Reachability is harder for directed than for undirected finite graphs, J. Symbolic Logic 55, 113-150.
- [AF94] Arora, S., and Fagin, R. (1994), On winning strategies in Ehrenfeucht-Fraïssé games, IBM Research Report RJ9833, June 1994.
- [Ajt83] Ajtai, M. (1983), Σ_1^1 -formulae on finite structures, Annals of Pure and Applied Logic 24, 1-48.
- [AU79] Aho, A. V., and Ullman, J. D. (1979), Universality of data retrieval languages, *in* "Proceedings, 6th ACM Symposium on Principles of Programming Languages," pp. 110-117.
- [CH82] Chandra, A., and Harel, D. (1982), Structure and complexity of relational queries, J. Comput. System Sci. 25, 99–128.
- [Cha88] Chandra, A. (1988), Theory of database queries, in "Proceedings, 7th ACM Symposium on Principles of Database Systems," pp. 1–9.
- [Cos93] Cosmadakis, S. (1993), Logical reducibility and monadic NP, in "Proceedings, 33rd IEEE Symposium on Foundations of Computer Science," pp. 52-61.
- [Cou90] Courcelle, B. (1990), The monadic second-order logic of graphs I: recognizable sets of finite graphs, *Inform. and Comput.* 85, 12-75.
- [dR87] de Rougemont, M. (1987), Second-order and inductive definability on finite structures, Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 33, 47-63.
- [Ehr61] Ehrenfeucht, A. (1961), An application of games to the completeness problem for formalized theories, *Fund. Math.* 49, 129–141.
- [End72] Enderton, H. B. (1972), "A Mathematical Introduction to Logic," Academic Press, New York.
- [Fag74] Fagin, R. (1974), Generalized first-order spectra and polynomial-time recognizable sets, in "Complexity of Computation, SIAM-AMS Proceedings, Vol. 7" (R. M. Karp, Ed.), pp. 43–73.
- [Fag75a] Fagin, R. (1975), Monadic generalized spectra, Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 21, 89-96.
- [Fag75b] Fagin, R. (1975), A spectrum hierarchy, Zeitschrift für Mathematische Logik und Grundlagen der Mathematik 21, 123–134.
- [Fag93] Fagin, R. (1993), Finite-model theory -a personal perspective, *Theoret. Comput. Sci.* 116, 3-31.
- [Fra54] Fraïssé, R. (1954), Sur quelques classifications des systèmes de relations, *Publ. Sci. Univ. Alger. Sér. A* 1, 35–182.

- [FSS84] Furst, M., Saxe, J. B., and Sipser, M. (1984), Parity, circuits, and the polynomial-time hierarchy, *Math. Systems Theory* 17, 13-27.
- [Gai82] Gaifman, H. (1982), On local and nonlocal properties, in "Logic Colloquium '81" (J. Stern, Ed.), pp. 105–135, North Holland, Amsterdam.
- [GJ79] Garey, M. R., and Johnson, D. S. (1979), "Computers and Intractability - A Guide to the Theory of NP-Completeness,"
 W. H. Freeman and Co., San Francisco.
- [Háj75] Hájek, P. (1975), On logics of discovery, in "Proceedings, 4th Conference on Mathematical Foundations of Computer Science, Lecture Notes in Computer Science 32," pp. 30–45, Springer-Verlag, Berlin.
- [Han65] Hanf, W. (1965), Model-theoretic methods in the study of elementary logic, in "The Theory of Models" (J. Addison, L. Henkin, and A. Tarski, Eds.), pp. 132-145, North Holland, Amsterdam.
- [Imm86] Immerman, N. (1986), Relational queries computable in polynomial time, *Inform. and Control* 68, 86-104.
- [Imm87] Immerman, N. (1987), Languages that capture complexity classes, SIAM J. Comput. 16, 760-778.
- [Imm89] Immerman, N. (1989), Descriptive and computational complexity, in "Computational Complexity Theory, Proceedings of Symposia in Applied Math., Vol. 38" (J. Hartmanis, Ed.), pp. 75-91, American Mathematical Society, Providence, RI.

- [JS74] Jones, N. G., and Selman, A. L. (1974), Turing machines and the spectra of first-order formulas, J. Symbolic Logic **39**, 139–150.
- [LY80] Lewis, J. M., and Yannakakis, M. (1980), The node deletion problem for hereditary properties is NP-complete, J. Comput. System Sci. 20, 219–230.
- [Sch94a] Schwentick, T. (1994), Graph connectivity and monadic NP, Research Report, Universität Mainz, August 1994.
- [Sch94b] Schwentick, T. (1994), Graph connectivity and monadic NP, in "Proceedings, 35th IEEE Symposium on Foundations of Computer Science," pp. 614-622.
- [Sho67] Shoenfield, J. R. (1967), "Mathematical Logic," Addison-Wesley, Reading, MA.
- [Spe91] Spencer, J. (1991), Threshold spectra via the Ehrenfeucht game, Discrete Appl. Math. 30, 235–252.
- [Sto77] Stockmeyer, L. J. (1977), The polynomial-time hierarchy, Theoret. Comput. Sci. 3, 1-22.
- [Tho91] Thomas, W. (1991), On logics, tilings, and automata, in "Proceedings, 18th International Colloquium on Automata, Languages, and Programming, Lecture Notes in Computer Science 510," pp. 441–454, Springer-Verlag, Berlin.
- [Tur84] Turán, G. (1984), On the definability of properties of finite graphs, Discrete Math. 49, 291-302.
- [Var82] Vardi, M. Y. (1982), The complexity of relational query languages, in "Proceedings, 14th ACM Symposium on Theory of Computing," pp. 137–146.