Math. Log. Quart. 43 (1997) 431 – 455

Mathematical Logic
Quarterly

Comparing the Power of Games on Graphs

Ronald Fagin[1)]

IBM Almaden Research Center, 650 Harry Road,
San Jose, California 95120-6099, U.S.A.[2)]

Abstract. The descriptive complexity of a problem is the complexity of describing the problem in some logical formalism. One of the few techniques for proving separation results in descriptive complexity is to make use of games on graphs played between two players, called the spoiler and the duplicator. There are two types of these games, which differ in the order in which the spoiler and duplicator make various moves. In one of these games, the rules seem to be tilted towards favoring the duplicator. These seemingly more favorable rules make it easier to prove separation results, since separation results are proven by showing that the duplicator has a winning strategy. In this paper, the relationship between these games is investigated. It is shown that in one sense, the two games are equivalent. Specifically, each family of graphs used in one game (the game with the seemingly more favorable rules for the duplicator) to prove a separation result can in principle be used in the other game to prove the same result. This answers an open question of Ajtai and the author from 1989. It is also shown that in another sense, the games are not equivalent, in that there are situations where the spoiler requires strictly more resources to win one game than the other game. This makes formal the informal statement that one game is easier for the duplicator to win.

Mathematics Subject Classification: 03D15, 03B15, 03C13, 68Q15.

Keywords: Descriptive complexity, Computational complexity, Ehrenfeucht-Fraïssé game, Ajtai-Fagin game, Finite model theory, Monadic NP.

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 (see [9]). There is an intimate connection between the descriptive complexity and the computational complexity. In particular (see [5]), 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. A consequence of this result is that NP=co-NP if and

[1)] A preliminary version of this paper appeared under the title "Comparing the power of monadic NP games" in Logic and Computational Complexity (D. LEIVANT, ed.), Springer Lecture Notes in Computer Science 960 (1995), pp. 414 – 425. The author is grateful to NEIL IMMERMAN, PHOKION KOLAITIS and MOSHE VARDI for helpful discussions. The author is also grateful to THOMAS WILKE for asking questions that led the author to Theorems 5.3 and 5.5.

[2)] e-mail: fagin@almaden.ibm.com; http://www.almaden.ibm.com/cs/people/fagin/

only if existential and universal second-order logic have the same expressive power over finite structures, i.e., if and only if $\Sigma_1^1 = \Pi_1^1$.

One way of attacking these difficult questions is to restrict the classes under consideration.Instead of considering Σ_1^1 (=NP) and Π_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). Following FAGIN, STOCKMEYER and VARDI [8], we refer to the restricted classes as *monadic NP* (resp., *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.

As a first step in this program, the author [6] separated monadic NP from monadic co-NP. Specifically, it was shown that connectivity of finite graphs is not in monadic NP, although it is easy to see that it is in monadic co-NP. The proof that connectivity is not in monadic NP makes use of a certain type of Ehrenfeucht-Fraïssé game on graphs played between two players, called the *spoiler* and the *duplicator*. The game involves coloring steps (where the players color nodes of the graphs) and selection steps (where the players select nodes of the graphs, round by round). We call this game the (original) *monadic NP game*. In this game, the duplicator selects two graphs G_0 and G_1, where G_0 is connected and G_1 is not. The spoiler then colors G_0, and the duplicator colors G_1. They then play a first-order Ehrenfeucht-Fraïssé game on these colored graphs, where, as usual, the spoiler tries to expose differences in the graphs, and the duplicator tries to cover up these differences. A necessary and sufficient condition for proving that connectivity is not in monadic NP is to show that for each choice of parameters (number of colors and number of first-order rounds), there are graphs G_0 and G_1 on which the duplicator has a winning strategy. By showing that, indeed, the duplicator has a winning strategy, the author showed that connectivity is not in monadic NP.

Later, AJTAI and the author [1] continued this program by showing that (s,t)-connectivity of directed graphs (otherwise known as *directed reachability*) is not in monadic NP. They made use of a modified game, which is now often referred to as the *Ajtai-Fagin monadic* NP *game*. Here the duplicator selects a graph G_0 that is (s,t)-connected, and the spoiler colors G_0. Then the duplicator selects and colors a graph G_1 that is not (s,t)-connected. The game again concludes with a first-order game. The difference between the Ajtai-Fagin game and the original game is that in the Ajtai-Fagin game, the spoiler must commit himself to a coloring of G_0 before seeing G_1. Putting it another way, the duplicator can wait to decide on his choice of G_1 until he sees how the spoiler colors G_0. Because the change in rules between the original game and the Ajtai-Fagin game favors the duplicator, on the face of it the Ajtai-Fagin game is "easier for the duplicator to win", which makes it easier to prove that the duplicator has a winning strategy. In fact, AJTAI and the author introduced their variation on the original game because they did not see how to prove that the duplicator has a winning strategy in the original game. However, they were able to prove that the duplicator has a winning strategy in their variation of the game. Since the duplicator has a winning strategy, directed reachability is not in monadic NP.

There is some mystery about the relationship between the Ajtai-Fagin game and the original game. On the one hand, the two games are equivalent, in the sense that in both cases, the existence of a winning strategy for the spoiler is a necessary and sufficient condition for a class to be in monadic NP. Thus, in both cases, showing that a problem is not in monadic NP corresponds precisely to showing that the duplicator has a winning strategy. On the other hand, as we noted, the Ajtai-Fagin game seems intuitively to be easier for the duplicator to win. Because of the fundamental role of Ehrenfeucht-Fraïssé games as tools in descriptive complexity, it is important to understand better the difference in power of the Ajtai-Fagin game and the original game. In this paper, we explore this difference.

For the sake of generality, we consider not just monadic NP games, but the more general (δ-ary) NP games, where the players color not just points, but δ-tuples of points for some fixed positive integer δ (the monadic case corresponds to $\delta = 1$). Similarly to before, the difference between the NP game and the Ajtai-Fagin NP game is that in the Ajtai-Fagin NP game, the spoiler must commit himself to a coloring of δ-tuples of points of G_0 before seeing G_1. Let us define *δ-ary NP* to be the restriction where second-order quantification is allowed only over δ-ary relations. Again, monadic NP is the case where $\delta = 1$. Binary NP (which corresponds to $\delta = 2$) is studied in [3]. The full class NP (that is, Σ_1^1) is the union over δ of δ-ary NP.

Let δ be fixed. In both the NP game and the Ajtai-Fagin NP game, a class C is given (such as the class of connected graphs). Then various graphs are selected and δ-tuples are colored by the players, and a first-order game is played on these colored graphs. The equivalence of the games corresponds to the fact that for each class C, the duplicator has a winning strategy in the original game for each choice of the remaining parameters (number of colors and number of rounds) if and only if the duplicator has a winning strategy in the Ajtai-Fagin game for each choice of the remaining parameters.

In this paper, we investigate the relationship between the original game and the Ajtai-Fagin game at a finer level. We show that in one sense, even at a finer level, Ajtai-Fagin NP games are no stronger than the original NP games. This sense corresponds to fact that in a game-theoretic proof that a class is not in δ-ary NP, the same families of graphs can be used in the original game as in the Ajtai-Fagin game. We also show that in another sense, Ajtai-Fagin games are stronger, in that there are situations where the spoiler requires more resources (colors) to win the Ajtai-Fagin game than the original game, when the choices of graphs are fixed. We now explain the details a little more.

In a game-theoretic proof that a specific class of graphs is not in δ-ary NP, the duplicator inevitably restricts himself to selecting graphs only of a certain type. For example, in the proof that connectivity is not in monadic NP (see [6]), the graph G_0 is a cycle, and G_1 is a disjoint union of two cycles. In the proof that directed reachability is not in monadic NP (see [1]), the graph G_0 is a path from s to t along with certain backedges, and G_1 is the result of deleting one forward edge from G_0. We show that each family of graphs used in the Ajtai-Fagin game to prove that a problem is not in δ-ary NP can in principle be used in the original game to prove the same result (where for a given choice of parameters, bigger graphs of the same type are used for the original game than for the Ajtai-Fagin game). For example, in the case of showing that directed reachability is not in monadic NP, we prove that for every choice of number of colors and number of rounds, the duplicator has a winning strategy in the (original) monadic NP game where the graph G_0 is a path from s to t along with certain backedges, and G_1 is the result of deleting one forward edge from G_0. This answers an open question of AJTAI and the author [1].

How do we obtain this result? First, we generalize the framework of the games. Rather than saying that the duplicator selects G_0 from a class $\mathcal{C}$, and selects G_1 from the complement $\overline{\mathcal{C}}$, we instead consider a more general game, where the duplicator selects G_0 from a class $\mathcal{G}_0$, and G_1 from a class $\mathcal{G}_1$. Intuitively, $\mathcal{G}_0$ and $\mathcal{G}_1$ correspond to the classes of graphs that are actually used in the games. For example, in the case of proving that connectivity is not in monadic NP, the class $\mathcal{G}_0$ would contain only graphs that are cycles, and the class $\mathcal{G}_1$ would contain only graphs that are the disjoint union of two cycles. There are once again two versions, one corresponding to the original game, and one to the Ajtai-Fagin game. In the first version of this new game, the duplicator selects G_1 before the spoiler has colored the δ-tuples of points of G_0; in the Ajtai-Fagin version, the duplicator selects G_1 after the spoiler has colored the δ-tuples of points of G_0. We show that given δ, for each choice of the number c of colors and the number r of rounds there are c' and r' such that for every choice of $\mathcal{G}_0$ and $\mathcal{G}_1$ where the duplicator has a winning strategy in the Ajtai-Fagin version of the new game with parameters c' and r', the duplicator also has a winning strategy in the first version of this game with parameters c and r (in fact, we can take $r' = r$). This result tells us that the same families of graphs can be used in the original game as in the Ajtai-Fagin game (such as to prove that a class is not in monadic NP). Intuitively, for a given choice of c, r, we use bigger graphs in the original game than in the Ajtai-Fagin game, since in the original game we use $\mathcal{G}_0$, $\mathcal{G}_1$ that correspond to the Ajtai-Fagin game with more colors (since $c' \geq c$).

We now consider a sense in which Ajtai-Fagin games are stronger. Here, we investigate the resources involved in the games. Specifically, we consider the number of colors required for the spoiler to win when the choices of graphs are fixed. Since the spoiler is trying to expose differences between G_0 and G_1, and the duplicator is trying to cover up these differences, it helps the spoiler for there to be more colors. We show that there are situations even when $\delta = 1$ where the spoiler requires strictly more colors to win the Ajtai-Fagin game than the original game. Thus, in such situations, it is indeed true, in a precise sense, that it is easier for the duplicator to win the Ajtai-Fagin game than the original game.

Our analysis gives a nonelementary upper bound on the number of extra colors that are required for the spoiler to win the Ajtai-Fagin game than the original game. We conjecture that there is also a nonelementary lower bound.

In Section 2, we give some definitions and conventions. In Section 3, we discuss Ehrenfeucht-Fraïssé games (see [4, 7]). In particular, we define both the original δ-ary NP game and the Ajtai-Fagin δ-ary NP game. In Section 4, we state and prove a useful result that is implicit in [6]. In Section 5, we introduce a notion of ***inseparability***, with which we can precisely define the notion of "the graphs used in a game". In Section 6, we demonstrate a strong sense in which the original game and the Ajtai-Fagin game are equivalent. In particular, we show that the same graphs used in the Ajtai-Fagin game to prove that a problem is not in δ-ary NP can in principal be used in the original game (for a different choice of parameters c and r) to prove this result. In Section 7, we discuss this strong equivalence in the context of AJTAI and the author's proof [1] that directed reachability is not in monadic NP. In Section 8, we show that the Ajtai-Fagin game is stronger than the original game, in that there are situations where the spoiler requires strictly more colors to win the Ajtai-Fagin game than the original game. In Section 9, we consider how many more colors may be necessary. In Section 10, we give our conclusions, and state some open problems.

2 Definitions and conventions

We begin with a few conventions. For convenience, we shall usually discuss only *graphs* (usually *directed* graphs, sometimes with distinguished points s and t), but everything we say can be generalized to arbitrary structures. We are also interested in "colored graphs", where each δ-tuple of points has some color. We assume throughout this paper that we are restricting our attention to *finite* graphs (and so are doing finite model theory), although all of the results hold also without this assumption. If G is a structure and φ is a sentence, then we use the usual Tarskian truth semantics to define what it means for φ to be *true* or *satisfied* in G, written $G \models \varphi$.

A Σ_1^1 *sentence* is a sentence of the form $\exists A_1 ... \exists A_k \, \psi$, where ψ is first-order and where the A_i's are relation symbols. We assume for convenience that the relation symbols A_i are all of the same arity δ (it is easy to see that this gives us no loss in expressive power). We refer to ψ as the *first-order part* of $\exists A_1 ... \exists A_k \, \psi$. As an example, we now construct a Σ_1^1 sentence that says that a graph (with edge relation denoted by P) is 3-colorable. In this sentence, the three colors are represented by the monadic 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 points with the same color are connected by an edge". Thus, ψ_2 is

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

The following sentence, which is Σ_1^1, then says "The graph is 3-colorable":

(1) $\exists A_1 \exists A_2 \exists A_3 (\psi_1 \wedge \psi_2)$.

A Σ_1^1 sentence $\exists A_1 ... \exists A_k \, \psi$, where ψ is first-order, is said to be *δ-ary* if each of the A_i's is δ-ary, that is, the existential second-order quantifiers quantify only over δ-ary relations. When $\delta = 1$, so that the existential second-order quantifiers quantify only over sets, we say that the sentence is *monadic*. A class C of graphs is said to be (*δ-ary*) Σ_1^1 if it is the class of all graphs that obey some fixed (δ-ary) Σ_1^1 sentence. One reason that Σ_1^1 classes are of great interest is the result [5] that the collection of Σ_1^1 classes coincides with the complexity class NP. For this reason, as we noted earlier, we follow FAGIN, STOCKMEYER, and VARDI [8] by referring to the collection of monadic Σ_1^1 classes as *monadic NP*, and more generally, to δ-ary Σ_1^1 classes as *δ-ary NP*. We often refer to a class of graphs by a defining property, for example, 3-colorability. As we saw above, 3-colorability is in monadic NP.

In a δ-ary Σ_1^1 sentence $\exists A_1 \ldots \exists A_k\, \psi$, the A_i's are, intuitively, allowed to intersect in arbitrary ways. We shall sometimes find it convenient to use slightly more restricted quantifiers, which we call *color quantifiers*. For a fixed choice of δ, we write $\exists\{B_1, \ldots, B_c\}\, \psi$ to mean, intuitively, that there is a way to color the δ-tuples with c colors such that ψ holds; we think of $\exists\{B_1, \ldots, B_c\}$ as a δ-ary color quantifier. More formally, we define $\exists\{B_1, \ldots, B_c\}\, \psi$ to be an abbreviation for $\exists A_1 \ldots \exists A_c\, (\psi' \wedge \varrho)$, where $A_1, \ldots, A_c$ are δ-ary relation symbols that do not appear in ψ, where ψ' is the result of replacing B_i by A_i for each i with $1 \leq i \leq c$, and where ϱ is the sentence that is the obvious generalization of sentence ψ_1 above from 3 colors to c colors, except that x represents an arbitrary δ-tuple, rather than a point. As a natural example of the use of color quantifiers, let ψ_2' be the first-order sentence obtained from ψ_2 above by replacing each A_i by the corresponding (unary) B_i, for $i = 1, 2, 3$. Then the Σ_1^1 sentence (1) above, which says "The graph is 3-colorable", is equivalent to the sentence $\exists\{B_1, B_2, B_3\}\, \psi_2'$. Note that, in a natural way, the color quantifier $\exists\{B_1, \ldots, B_c\}$, which involves c colors in the color quantifier, can be simulated by $\exists A_1 \ldots \exists A_{\lceil \log c \rceil}$, which involves $\lceil \log c \rceil$ quantifiers. In particular, when $\exists\{B_1, \ldots, B_c\}$ is a δ-ary color quantifier, then every sentence $\exists\{B_1, \ldots, B_c\}\, \psi$ where ψ is first-order is δ-ary Σ_1^1.

The *quantifier depth* $QD(\varphi)$ of a first-order formula φ is defined recursively as follows.

- $QD(\varphi) = 0$ if φ is quantifier-free,
- $QD(\neg\varphi) = QD(\varphi)$, $\quad QD(\varphi_1 \wedge \varphi_2) = \max\{QD(\varphi_1), QD(\varphi_2)\}$,
- $QD(\exists x\, \varphi) = 1 + QD(\varphi)$.

3 Ehrenfeucht-Fraïssé games

Among the few tools of model theory that "survive" when we restrict our attention to finite structures are Ehrenfeucht-Fraïssé games [4, 7]. 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, G_0 and G_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 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 structures, and the duplicator selects a point in the other structure. 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 substructure of G_0 induced by $\langle a_1, \ldots, a_r \rangle$ is isomorphic to the substructure of G_1 induced by $\langle b_1, \ldots, b_r \rangle$, under the function that maps a_i onto b_i for $1 \leq i \leq r$. For example, in the case where G_0 and G_1 are colored graphs (where each δ-tuple of points is colored with one of c possible colors), then for the duplicator to win, (a) $a_i = a_j$ iff $b_i = b_j$, for each i, j;

(b) $\langle a_i, a_j \rangle$ is an edge in G_0 iff $\langle b_i, b_j \rangle$ is an edge in G_1, for each i, j; and (c) $\langle a_{i_1}, \ldots, a_{i_\delta} \rangle$ has the same color as $\langle b_{i_1}, \ldots, b_{i_\delta} \rangle$, for each $i_1, \ldots, i_\delta$. Otherwise, the spoiler wins.[3)] 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 [4, 7]) shows why these games are of interest. If $\mathcal{C}$ is a class of structures, then let $\overline{\mathcal{C}}$ be the complement of $\mathcal{C}$, that is, the class of structures not in $\mathcal{C}$.[4)]

Theorem 3.1. *Let $\mathcal{C}$ be a class of structures. $\mathcal{C}$ is first-order definable iff there is r such that whenever $G_0 \in \mathcal{C}$ and $G_1 \in \overline{\mathcal{C}}$, 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, δ-ary NP game, and which we shall call a *(δ, c, r)-game* for short. This game was essentially introduced in [6] 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 δ-tuples of points of G_0, using the colors in C, and then the duplicator colors each of the δ-tuples of 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 δ-tuples of points of G_0, not G_1. The game then concludes with an r-game. The duplicator now wins if the substructure of G_0 induced by $\langle a_1, \ldots, a_r \rangle$ is isomorphic to the substructure of G_1 induced by $\langle b_1, \ldots, b_r \rangle$, under the function that maps a_i onto b_i for $1 \le i \le r$. When $\delta = 1$ (where individual points are colored), we get the *monadic* NP *game*, which we shall sometimes refer to as the *(c, r)-game*.

The following theorem (essentially from [6]) is analogous to Theorem 3.1.

Theorem 3.2. *Let $\mathcal{C}$ be a class of graphs, and let δ be a positive integer. $\mathcal{C}$ is in δ-ary* NP *iff there are c, r such that whenever $G_0 \in \mathcal{C}$ and $G_1 \in \overline{\mathcal{C}}$, then the spoiler has a winning strategy in the (δ, c, r)-game over G_0, G_1.*

In [6] 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 the author [1] found it convenient, for reasons we shall see shortly, to consider games over a class $\mathcal{C}$ of graphs. The rules of an r-game over $\mathcal{C}$ are as follows. The duplicator

---

3) If we are dealing with structures that have constant symbols $c_1, \ldots, c_z$, then we slightly modify the definition of when the duplicator wins. Specifically, let a_{r+i} denote the interpretation in G_0 of c_i, and let b_{r+i} denote the interpretation in G_1 of c_i, for $1 \le i \le z$. The duplicator wins if the substructure of G_0 induced by $\langle a_1, \ldots, a_{r+z} \rangle$ is isomorphic to the substructure of G_1 induced by $\langle b_1, \ldots, b_{r+z} \rangle$, under the function that maps a_i onto b_i for $1 \le i \le r + z$. Thus, we consider the substructure that is generated by not only the points selected in the structures, but also by the interpretations of the constant symbols. This modification is necessary for Theorem 5.1 to hold in the presence of constant symbols.

4) Of course, we are restricting attention to structures of a given similarity type. For example, if the similarity type is that of graphs, then $\overline{\mathcal{C}}$ contains the graphs not in $\mathcal{C}$.

begins by selecting a member of $\mathcal{C}$ to be G_0, and a member of $\overline{\mathcal{C}}$ 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 (δ, c, r)-game over $\mathcal{C}$. The rules are as follows.

1. The duplicator selects a member of $\mathcal{C}$ to be G_0.
2. The duplicator selects a member of $\overline{\mathcal{C}}$ to be G_1.
3. The spoiler colors the δ-tuples of points of G_0 with the c colors.
4. The duplicator colors the δ-tuples of points of 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. *Let $\mathcal{C}$ be a class of graphs.*

(a) *$\mathcal{C}$ is first-order definable iff there is r such that the spoiler has a winning strategy in the r-game over $\mathcal{C}$.*

(b) *$\mathcal{C}$ is in δ-ary* NP *iff there are c, r such that the spoiler has a winning strategy in the (δ, c, r)-game over $\mathcal{C}$.*

Theorem 3.3 says that r-games are sound and complete for proving that a class is not first-order, and that (δ, c, r)-games are sound and complete for proving that a class is not in δ-ary NP. Thus, in the case of δ-ary NP, we have the following.

Soundness: To show that $\mathcal{C}$ is not in δ-ary NP, it is sufficient to show that for every c, r, the duplicator has a winning strategy in the (δ, c, r)-game over $\mathcal{C}$.

Completeness: If $\mathcal{C}$ is not in δ-ary NP, then in principle this can be shown by a game argument (that is, for every c, r, the duplicator has a winning strategy in the (δ, c, r)-game over $\mathcal{C}$).

We now explain why AJTAI and the author allowed 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 the author wished to prove that directed (s,t)-connectivity (also known as directed reachability) 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 reachability is not in monadic NP). Their new game, which is usually called the *Ajtai-Fagin (c, r)-game* (or, in our case, the Ajtai-Fagin (δ, 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 (δ, c, r)-game by reversing the order of the second and third moves. Thus, the rules of the Ajtai-Fagin (δ, c, r)-game are as follows.

1. The duplicator selects a member of $\mathcal{C}$ to be G_0.
2. The spoiler colors the δ-tuples of points of G_0 with the c colors.
3. The duplicator selects a member of $\overline{\mathcal{C}}$ to be G_1.
4. The duplicator colors the δ-tuples of points of 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 (δ, c, r)-game, the spoiler must commit himself to a coloring of G_0 before knowing what G_1 is. In order to contrast it with the Ajtai-Fagin (δ, c, r)-game, we may sometimes refer to the (δ, c, r)-game as the *original* (δ, c, r)-game (or the *original* δ-ary NP game). In spite of the fact that it seems to be harder for the spoiler to win the Ajtai-Fagin (δ, c, r)-game than the original (δ, c, r)-game, we have the following analogue (which is a slight generalization of a theorem of [1]) to Theorem 3.3(b).

Theorem 3.4. *Let C be a class of graphs, and let δ be a positive integer. The class C is in δ-ary* NP *iff there are c, r such that the spoiler has a winning strategy in the Ajtai-Fagin (δ, c, r)-game over C.*

Thus, in the same sense as before, Theorem 3.4 says that Ajtai-Fagin (δ, c, r)-games are sound and complete for proving that a class is not in δ-ary NP.

The next theorem is an immediate consequence of Theorems 3.3(b) and 3.4:

Theorem 3.5. *Let C be a class of graphs, and let δ be a positive integer. The following are equivalent.*

(i) *For every c, r, the duplicator has a winning strategy in the original (δ, c, r)-game over C.*

(ii) *For every c', r', the duplicator has a winning strategy in the Ajtai-Fagin (δ, c', r')-game over C.*

Theorem 3.5 gives a precise sense in which the original δ-ary NP game and the Ajtai-Fagin δ-ary NP game are equivalent. Later, we shall see stronger versions of this equivalence.

4 A game-theoretic tool

In this section, we state and prove a useful result that is implicit in [6]. The derivations follow those in [6].

Let us define an *r-sentence* to be a first-order sentence of quantifier depth at most r, and a *(δ, c, r)-sentence* to be a sentence of the form $\exists\{B_1, \ldots, B_c\}\,\psi$, where ψ is an r-sentence, and where $\exists\{B_1, \ldots, B_c\}$ is a δ-ary color quantifier.

Theorem 4.1. *For each graph G_0 and each triple δ, c, r of positive integers, there is a sentence $\sigma(G_0, \delta, c, r)$ that is a conjunction of (δ, c, r)-sentences, such that for each graph G_1, the duplicator has a winning strategy in the (δ, c, r)-game over G_0, G_1 iff $G_1 \models \sigma(G_0, \delta, c, r)$. For each δ, c, r, there are only a finite number of distinct inequivalent sentences $\sigma(G_0, \delta, c, r)$.*

Proof. Let δ, c, r be fixed. Let P be a binary relation symbol that represent the graph relation, and let $B_1, \ldots, B_c$ represent the c colors. The *underlying graph* of the colored graph is obtained by ignoring the coloring (that is, taking only the graph relation).

For each integer m with $0 \leq m \leq r$, we define an *m-type*, by induction on m. Let $v_1, \ldots, v_r$ be distinct individual variables, and define the *atomic formulas over* $\{v_1, \ldots, v_r\}$ to be all formulas of one of the following forms:

1. $v_i = v_j$, where $1 \leq i \leq r$ and $1 \leq j \leq r$;

2. Pv_iv_j, where $1 \leq i \leq r$ and $1 \leq j \leq r$;
3. $B_\ell v_{i_1} \ldots v_{i_\delta}$, where $1 \leq \ell \leq c$, and $1 \leq i_j \leq r$ for $1 \leq j \leq \delta$.

Let $\alpha_1, \ldots, \alpha_s$ be the distinct atomic formulas over $\{v_1, \ldots, v_r\}$. We start the induction by defining a 0-type to be a formula of the form $\beta_1 \wedge \cdots \wedge \beta_s$, where each β_i is either α_i or $\neg\alpha_i$. Intuitively, a 0-type is a complete description of how the variables $v_1, \ldots, v_r$ relate to each other (in terms of which are equal, and which have edges between them in the graph relation) and which colors the δ-tuples of them have. We have not bothered to require that each 0-type be consistent: for example, a 0-type for $\delta = 1$ could have as conjuncts each of the formulas $v_1 = v_2$, B_1v_1, and $\neg B_1v_2$. Of course, it would be straightforward to make such a restriction. Note that a 0-type is of quantifier depth 0.

For each set A of m-types (where $0 \leq m < r$), we inductively define the following formula to be an $(m+1)$-type:

$$\bigwedge\{\exists v_{r-m}\varphi : \varphi \in A\} \wedge \bigwedge\{\forall v_{r-m}\neg\varphi : \varphi \notin A\}.$$

Assume that H is a colored graph, that $a_1, \ldots, a_m$ are points of H, and that φ is a formula with free variables $v_1, \ldots, v_m$. We write $H \models \varphi[v_1 \cdots v_m \mid a_1 \cdots a_m]$ to mean that φ is satisfied in H when v_i is interpreted by a_i, for $1 \leq i \leq m$. It is easy to prove by induction on m that if $0 \leq m \leq r$, then

1. each m-type is of quantifier depth m;
2. each m-type has as free variables precisely $v_1, \ldots, v_{r-m}$;
3. there is exactly one m-type φ such that $H \models \varphi[v_1 \cdots v_{r-m} \mid a_1 \cdots a_{r-m}]$;
4. there are only finitely many distinct m-types.

It follows from the inductive assumptions (when $m = r$) that each r-type is of quantifier depth r, and the number of r-types is finite, where the total number depends only on δ, c and r. Let us denote by $\tau(H, r)$ the r-type φ such that $H \models \varphi$. It is easy to see that if H and H' are colored graphs, then $H \sim_r H'$ iff $H \models \tau(H, r)$: the duplicator's winning strategy is to make sure that after the mth move, if $a_1, \ldots, a_m$ (respectively, $b_1, \ldots, b_m$) are the points that have been picked in H (respectively, H'), then

$$H \models \psi[v_1 \cdots v_m \mid a_1 \cdots a_m] \text{ iff } H' \models \psi[v_1 \cdots v_m \mid b_1 \cdots b_m]$$

for the same $(r-m)$-type ψ.

Let G_0 be a graph. Let $\mathcal{H}$ be the set of all colored graphs H where the underlying graph of H is G_0. Since the number of r-types is finite and depends only on δ, c and r, there is a set $\mathcal{G} \subseteq \mathcal{H}$ such that the r-type of each member of $\mathcal{H}$ is the same as the r-type of some member of $\mathcal{G}$, and such that the cardinality of $\mathcal{G}$ is finite and depends only on δ, c and r. (Note that the set $\mathcal{H}$ itself is finite, since there are only a finite number of ways to color the δ-tuples of G_0 with c colors. However, the cardinality depends on the size of G_0, and we need an upper bound that depends only on δ, c and r.) Define $\sigma(G_0, \delta, c, r)$ to be the conjunction over $\mathcal{G}$ of all (δ, c, r)-sentences $\exists\{B_1, \ldots, B_c\}\tau(H, r)$, where $H \in \mathcal{G}$. It is easy to see from what we have said that there are only a finite number of distinct inequivalent sentences of the form $\sigma(G_0, \delta, c, r)$, and this number depends only on δ, c and r. It is straightforward to verify that for each graph G_1, the duplicator has a winning strategy in the (δ, c, r)-game over G_0, G_1 iff $G_1 \models \sigma(G_0, \delta, c, r)$. This proves the theorem. □

The next theorem, that we shall find useful later, uses the notion of r-type as defined in the proof of Theorem 4.1.

Theorem 4.2. *Each sentence* $\sigma(G_0, \delta, c, r)$ *is of the form* $\bigwedge_{i \in I} \exists\{B_1, \dots, B_c\}\psi_i$, *where each* ψ_i *is an* r*-type.*

5 Inseparability

In this section, we introduce a notion of *inseparability*. In this section and the next section, we use inseparability to give stronger versions of the equivalence between the original δ-ary NP game and the Ajtai-Fagin δ-ary NP game. This notion of inseparability allows us to make sense of the notion of "the graphs used in a game", so that we can consider statements such as "the same family of graphs used in the Ajtai-Fagin δ-ary NP game can be used in the original δ-ary NP game to prove that a problem is not in δ-ary NP".

We defined quantifier depth in Section 2. Quantifier depth is closely related to first-order games. We say that two structures G_0 and G_1 are *r-inseparable*, written $G_0 \sim_r G_1$, if the duplicator has a winning strategy in the r-game over G_0, G_1. If it is the spoiler who has a winning strategy, then we say that G_0 and G_1 are *r-separable*. Since the game is finite, and there are no ties, the spoiler has a winning strategy iff the duplicator does not; therefore, G_0 and G_1 are r-separable iff they are not r-inseparable.

We say that the sentence γ *separates* G_0 *from* G_1 if $G_0 \models \gamma$ and $G_1 \not\models \gamma$. The next theorem shows why r-separability is an important concept (and why we use the word "separability").

Theorem 5.1. *Let* r *be a positive integer. The following are equivalent.*

(i) G_0 *and* G_1 *are* r*-separable.*

(ii) *There is an* r*-sentence that separates* G_0 *from* G_1.

Proof. The proof that (i) $\Rightarrow$ (ii) follows as in the proof of Theorem 4.1. We can prove that (ii) $\Rightarrow$ (i) by a fairly straightforward induction on r. □

Let us now pass from the first-order case to the δ-ary Σ_1^1 case. Similarly to before, we say that the graphs G_0 and G_1 are (δ, c, r)*-inseparable* if the duplicator has a winning strategy in the (δ, c, r)-game over G_0, G_1. The next result is analogous to Theorem 5.1.

Theorem 5.2. *The following are equivalent.*

(i) G_0 *and* G_1 *are* (δ, c, r)*-separable.*

(ii) *There is a* (δ, c, r)*-sentence that separates* G_0 *from* G_1.

Proof.

(i) $\Rightarrow$ (ii). Let $\sigma(G_0, \delta, c, r)$ be as in Theorem 4.1. Assume that (i) holds. By Theorem 4.1, it follows immediately that $G_1 \not\models \sigma(G_0, \delta, c, r)$, and hence $G_1 \not\models \gamma$ for some conjunct γ of $\sigma(G_0, \delta, c, r)$. Also, Theorem 4.1 implies easily that $G_0 \models \sigma(G_0, \delta, c, r)$, and so $G_0 \models \gamma$. Since γ is a (δ, c, r)-sentence, this proves (ii).

(ii) $\Rightarrow$ (i). Let $\exists\{B_1, \ldots, B_c\}\,\psi$ be a (δ, c, r)-sentence such that

$$G_0 \models \exists\{B_1, \ldots, B_c\}\,\psi \quad \text{and} \quad G_1 \not\models \exists\{B_1, \ldots, B_c\}\,\psi.$$

Then there is a way to color the δ-tuples of G_0 with c colors to obtain a colored graph G_0' so that $G_0' \models \psi$, and no matter how the δ-tuples of G_1 are colored with c colors to obtain G_1', necessarily $G_1' \not\models \psi$. Since ψ is an r-sentence, it follows from Theorem 5.1 that G_0' and G_1' are r-separable. It follows easily that the spoiler has a winning strategy in the (δ, c, r)-game over G_0, G_1. Therefore, (i) holds. □

Although r-inseparability, which we considered in the first-order case, is an equivalence relation, (δ, c, r)-inseparability is not an equivalence relation in general. It is easy to see that (δ, c, r)-inseparability is transitive and reflexive. However, it is not necessarily symmetric. The fact that (δ, c, r)-inseparability is not symmetric is reflected in the fact that unlike the situation in Theorem 5.1, the second condition of Theorem 5.2 is not symmetric in the roles of G_0 and G_1. It is possible for the second condition of Theorem 5.2 to hold, and yet for the analogous condition where the roles of G_0 and G_1 are reversed to fail.

To make precise the notion of "the graphs used in the game", we need to consider a notion of separability of *classes* $\mathcal{G}_0$ and $\mathcal{G}_1$ of graphs. We begin by recalling the rules of the original (δ, c, r)-game over $\mathcal{C}$:

1. The duplicator selects a member of $\mathcal{C}$ to be G_0.
2. The duplicator selects a member of $\overline{\mathcal{C}}$ to be G_1.
3. The spoiler colors the δ-tuples of points of G_0 with the c colors.
4. The duplicator colors the δ-tuples of points of G_1 with the c colors.
5. The spoiler and duplicator play an r-game.

We now define a variation. Let $\mathcal{G}_0$ and $\mathcal{G}_1$ be classes of graphs, and let c, r be positive integers. We define the *original* (δ, c, r)-*game over* $\mathcal{G}_0, \mathcal{G}_1$ to have the following rules.

1. The duplicator selects a member of $\mathcal{G}_0$ to be G_0.
2. The duplicator selects a member of $\mathcal{G}_1$ to be G_1.
3. The spoiler colors the δ-tuples of points of G_0 with the c colors.
4. The duplicator colors the δ-tuples of points of G_1 with the c colors.
5. The spoiler and duplicator play an r-game.

The winner is decided as before. Intuitively, $\mathcal{G}_0$ and $\mathcal{G}_1$ correspond to the classes of graphs that are actually used in the games. For example, in the case of proving that connectivity is not in monadic NP, the class $\mathcal{G}_0$ would contain only graphs that are cycles, and the class $\mathcal{G}_1$ would contain only graphs that are the disjoint union of two cycles.

We say that $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-*inseparable* if the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. In particular, if the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{C}$, then $\mathcal{C}$, $\overline{\mathcal{C}}$ are (δ, c, r)-inseparable. As before, if it is the spoiler, rather than the duplicator, who has a winning strategy, then we say that $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-*separable*. Further, as before, if $\delta = 1$, then we refer simply to (c,r)-inseparability and (c,r)-separability. Note that if $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ,c,r)-inseparable, and $\mathcal{C}$ is an arbitrary class such that $\mathcal{G}_0 \subseteq \mathcal{C}$ and $\mathcal{G}_1 \subseteq \overline{\mathcal{C}}$, then by Theorem 3.3(b), it follows that $\mathcal{C}$ is not in δ-ary NP. In particular, the same $\mathcal{G}_0$, $\mathcal{G}_1$ can be used to prove simultaneously that many classes are not in δ-ary NP. For example, consider the fact that for every c, r, there exists $\mathcal{G}_0$ consisting only of graphs that are cycles, and $\mathcal{G}_1$ consisting only of graphs that are the disjoint union of two cycles, such that $\mathcal{G}_0$, $\mathcal{G}_1$ are (c,r)-inseparable (see [6]). This fact implies that if $\mathcal{C}$ is an arbitrary class that contains every graph that is a cycle but no graph that is the disjoint union of two cycles, then $\mathcal{C}$ is not in monadic NP. Examples of such classes $\mathcal{C}$ are the class of connected graphs and the class of Hamiltonian graphs.

Theorem 5.2 characterizes separability in the case of singleton sets $\mathcal{G}_0$ and $\mathcal{G}_1$. We now give a characterization that holds in the general case. We say that the sentence γ *separates* $\mathcal{G}_0$ *from* $\mathcal{G}_1$ if (a) $G_0 \models \gamma$ for every $G_0 \in \mathcal{G}_0$, and (b) $G_1 \not\models \gamma$ for every $G_1 \in \mathcal{G}_1$. A sentence is a *positive Boolean combination of* (δ,c,r)*-sentences* if it is in the smallest class that contains every (δ,c,r)-sentence and is closed under conjunction and disjunction. In the theorem below, $\sigma(G_0,\delta,c,r)$ is as in Theorem 4.1.

Theorem 5.3. *The following are equivalent.*

(i) $\mathcal{G}_0$ *and* $\mathcal{G}_1$ *are* (δ,c,r)*-separable.*

(ii) *There is a sentence that is a positive Boolean combination of* (δ,c,r)*-sentences that separates* $\mathcal{G}_0$ *from* $\mathcal{G}_1$.

(iii) *There is a sentence that is a disjunction of conjunctions of* (δ,c,r)*-sentences that separates* $\mathcal{G}_0$ *from* $\mathcal{G}_1$.

(iv) *The sentence* $\bigvee_{G_0\in\mathcal{G}_0}\sigma(G_0,\delta,c,r)$ *separates* $\mathcal{G}_0$ *from* $\mathcal{G}_1$.[5)]

Proof. We shall prove that (iv) $\Rightarrow$ (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i) $\Rightarrow$ (iv).

(iv) $\Rightarrow$ (iii). This is immediate, since the sentence $\bigvee_{G_0\in\mathcal{G}_0}\sigma(G_0,\delta,c,r)$ is a disjunction of conjunctions of (δ,c,r)-sentences.

(iii) $\Rightarrow$ (ii). This is immediate, since a disjunction of conjunctions is a positive Boolean combination.

(ii) $\Rightarrow$ (i). Let γ be a sentence that is a positive Boolean combination of (δ,c,r)-sentences that separates $\mathcal{G}_0$ from $\mathcal{G}_1$. Let G_0 be an arbitrary member of $\mathcal{G}_0$, and let G_1 be an arbitrary member of $\mathcal{G}_1$. Then $G_0 \models \gamma$ and $G_1 \not\models \gamma$. Since γ is a positive Boolean combination of (δ,c,r)-sentences, it is not hard to see that γ is equivalent to a sentence γ' of the form

$$\bigvee_{i_1\in I_1}\bigwedge_{i_2\in I_2}\bigvee_{i_3\in I_3}\cdots\bigwedge_{i_{k-1}\in I_{k-1}}\bigvee_{i_k\in I_k}\psi_{i_1,\ldots,i_k},$$

where each $\psi_{i_1,\ldots,i_k}$ is a (δ,c,r)-sentence. Since $G_0 \models \gamma'$, there is $i_1 \in I_1$ such that

$$G_0 \models \bigwedge_{i_2\in I_2}\bigvee_{i_3\in I_3}\cdots\bigwedge_{i_{k-1}\in I_{k-1}}\bigvee_{i_k\in I_k}\psi_{i_1,\ldots,i_k}. \tag{2}$$

Since $G_1 \not\models \gamma'$, it follows that

$$G_1 \not\models \bigwedge_{i_2\in I_2}\bigvee_{i_3\in I_3}\cdots\bigwedge_{i_{k-1}\in I_{k-1}}\bigvee_{i_k\in I_k}\psi_{i_1,\ldots,i_k}. \tag{3}$$

5) This is a finite disjunction, since by Theorem 4.1, for each fixed choice of δ, c, r, there are only a finite number of distinct inequivalent sentences of the form $\sigma(G_0,\delta,c,r)$.

Intuitively, (2) and (3) tell us that we can find some choice of $i_1 \in I_1$ so that we can "strip off the outer $\bigvee$".

From (3), we know that for some $i_2 \in I_2$, we have

(4) $\quad G_1 \not\models \bigvee_{i_3 \in I_3} \cdots \bigwedge_{i_{k-1} \in I_{k-1}} \bigvee_{i_k \in I_k} \psi_{i_1,\ldots,i_k}.$

From (2), it follows that

(5) $\quad G_0 \models \bigvee_{i_3 \in I_3} \cdots \bigwedge_{i_{k-1} \in I_{k-1}} \bigvee_{i_k \in I_k} \psi_{i_1,\ldots,i_k}.$

Intuitively, (4) and (5) tell us that we can find some choice of $i_1 \in I_1$ and $i_2 \in I_2$ so that we can "strip off the outer $\bigvee \bigwedge$".

By continuing this process a total of k times, we find that for some $i_1, \ldots, i_k$, the (δ, c, r)-sentence $\psi_{i_1,\ldots,i_k}$ separates G_0 from G_1. So from Theorem 5.2, it follows that G_0 and G_1 are (δ, c, r)-separable. That is, the spoiler has a winning strategy in the (δ, c, r)-game over G_0, G_1. Since G_0 is an arbitrary member of $\mathcal{G}_0$, and G_1 is an arbitrary member of $\mathcal{G}_1$, it follows that the spoiler has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. So $\mathcal{G}_0$ and $\mathcal{G}_1$ are (δ, c, r)-separable, as desired.

(i) $\Rightarrow$ (iv). Assume that $\mathcal{G}_0$ and $\mathcal{G}_1$ are (δ, c, r)-separable. Let γ be the formula $\bigvee_{G_0 \in \mathcal{G}_0} \sigma(G_0, \delta, c, r)$. Let G_0 be an arbitrary member of $\mathcal{G}_0$, and let G_1 be an arbitrary member of $\mathcal{G}_0$. We must show that $G_0 \models \gamma$ and $G_1 \not\models \gamma$. By Theorem 4.1, it follows immediately that $G_0 \models \sigma(G_0, \delta, c, r)$. Hence, $G_0 \models \gamma$. Assume now that $G_1 \models \gamma$; we shall derive a contradiction. Since $G_1 \models \gamma$, we have $G_1 \models \sigma(G_0^*, c, r)$ for some $G_0^* \in \mathcal{G}_0$. The duplicator therefore has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$, as we now show. The duplicator selects $G_0^* \in \mathcal{G}_0$ and $G_1 \in \mathcal{G}_1$. Since $G_1 \models \sigma(G_0^*, c, r)$, it follows from Theorem 4.1 that the duplicator has a winning strategy in the (δ, c, r)-game over G_0^*, G_1. So indeed, the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. This contradiction shows that $G_1 \not\models \gamma$, as desired. □

Let us now consider Ajtai-Fagin games. As before, let $\mathcal{G}_0$ and $\mathcal{G}_1$ be classes of graphs, and let δ, c, r be positive integers. We define the *Ajtai-Fagin (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$* to have the following rules.

1. The duplicator selects a member of $\mathcal{G}_0$ to be G_0.
2. The spoiler colors the δ-tuples of points of G_0 with the c colors.
3. The duplicator selects a member of $\mathcal{G}_1$ to be G_1.
4. The duplicator colors the δ-tuples of points of G_1 with the c colors.
5. The spoiler and duplicator play an r-game.

The winner is decided as before. As before, the difference between the original (δ, c, r)-game and Ajtai-Fagin (δ, c, r)-game is that in the Ajtai-Fagin game, the spoiler must commit himself to a coloring of G_0 before knowing which graph the duplicator selects as G_1.

In Section 6, we shall prove Theorem 5.4 below, which is a strengthened version of Theorem 3.5. In fact, we shall prove a result even stronger than Theorem 5.4.

Theorem 5.4. *Let $\mathcal{G}_0$, $\mathcal{G}_1$ be classes of graphs, and let δ be a positive integer. The following are equivalent.*

(i) *For every c, r, the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$.*

(ii) *For every c', r', the duplicator has a winning strategy in the Ajtai-Fagin (δ, c', r')-game over $\mathcal{G}_0$, $\mathcal{G}_1$.*

We say that $\mathcal{G}_0$, $\mathcal{G}_1$ are *Ajtai-Fagin (δ, c, r)-inseparable* (respectively *Ajtai-Fagin (δ, c, r)-separable*) if the duplicator (respectively the spoiler) has a winning strategy in the Ajtai-Fagin (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. As before, if $\delta = 1$, then we refer simply to Ajtai-Fagin (c, r)-inseparability and Ajtai-Fagin (c, r)-separability.

We now give a theorem that characterizes Ajtai-Fagin separability, analogous to the characterization of separability in Theorem 5.3.

Theorem 5.5. *The following are equivalent.*

(i) *$\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (δ, c, r)-separable.*

(ii) *There is a sentence that is a disjunction of (δ, c, r)-sentences that separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*

(iii) *For each $G_0 \in \mathcal{G}_0$, there is some (δ, c, r)-sentence γ_{G_0} that is a conjunct of $\sigma(G_0, \delta, c, r)$ such that the sentence $\bigvee_{G_0 \in \mathcal{G}_0} \gamma_{G_0}$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*[6)]

We shall prove that (iii) $\Rightarrow$ (ii) $\Rightarrow$ (i) $\Rightarrow$ (iii).

(iii) $\Rightarrow$ (ii). This is immediate, since γ_{G_0} is a (δ, c, r)-sentence.

(ii) $\Rightarrow$ (i). Let $\gamma_1, \ldots, \gamma_k$ be (δ, c, r)-sentences such that $\gamma_1 \vee \cdots \vee \gamma_k$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$. We must show that the spoiler has a winning strategy in the Ajtai-Fagin (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. Assume that the duplicator selects $G_0 \in \mathcal{G}_0$. Since $G_0 \models \gamma_1 \vee \cdots \vee \gamma_k$, we know that $G_0 \models \gamma_i$ for some i. Now γ_i is a (δ, c, r)-sentence $\exists\{B_1, \ldots, B_c\}\,\psi$. The spoiler now colors the δ-tuples of points of G_0 with the c colors so that the resulting colored graph G_0' satisfies ψ (this is possible, since $G_0 \models \exists\{B_1, \ldots, B_c\}\,\psi$). For whatever $G_1 \in \mathcal{G}_1$ that the duplicator now selects, we know that $G_1 \not\models \gamma_1 \vee \cdots \vee \gamma_k$. Hence, $G_1 \not\models \gamma_i$, that is, $G_1 \not\models \exists\{B_1, \ldots, B_c\}\,\psi$. Therefore, however the duplicator colors the δ-tuples of points of G_1, the resulting colored graph G_1' does not satisfy ψ. Since ψ is an r-sentence, it follows from Theorem 5.1 that G_0' and G_1' are r-separable. That is, the spoiler has a winning strategy in the r-game over G_0', G_1'. It follows easily from what we have shown that $\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (δ, c, r)-separable.

(i) $\Rightarrow$ (iii). Assume that $\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (δ, c, r)-separable. Let G_0 be an arbitrary member of $\mathcal{G}_0$. By Theorem 4.2, we know that $\sigma(G_0, \delta, c, r)$ is of the form $\bigwedge_{i \in I} \exists\{B_1, \ldots, B_c\}\,\psi_i$, where each ψ_i is an r-type. We see from the proof of Theorem 4.1 that each conjunct $\exists\{B_1, \ldots, B_c\}\,\psi_i$ of $\sigma(G_0, \delta, c, r)$ corresponds to a coloring of G_0. Let γ_{G_0} be the conjunct $\exists\{B_1, \ldots, B_c\}\,\psi_i$ of $\sigma(G_0, \delta, c, r)$ that corresponds to the coloring of G_0 in a winning strategy of the spoiler in the Ajtai-Fagin (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$ when the duplicator selects $G_0 \in \mathcal{G}_0$. Because of the

6) As in Theorem 5.3, this is a finite disjunction, since for each fixed choice of δ, c, r, there are only a finite number of distinct inequivalent conjuncts of sentences $\sigma(G_0, \delta, c, r)$.

choice of γ_{G_0}, it follows that for every choice of $G_1 \in \mathcal{G}_1$, and for every coloring by the duplicator of the δ-tuples of points of G_1 with c colors, the resulting colored graph does not satisfy ψ_i. Therefore, $G_1 \not\models \exists\{B_1, \ldots, B_c\}\,\psi_i$, that is, $G_1 \not\models \gamma_{G_0}$. Since G_0 is an arbitrary member of $\mathcal{G}_0$ and G_1 is an arbitrary member of $\mathcal{G}_1$, it follows that no member of $\mathcal{G}_1$ satisfies $\bigvee_{G_0 \in \mathcal{G}_0} \gamma_{G_0}$. Since $G_0 \models \gamma_{G_0}$, it follows that every member of $\mathcal{G}_0$ satisfies $\bigvee_{G_0 \in \mathcal{G}_0} \gamma_{G_0}$. Hence, $\bigvee_{G_0 \in \mathcal{G}_0} \gamma_{G_0}$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$, as desired. □

It is not hard to verify that for most purposes, we could have restricted our attention in this paper to cases where $\mathcal{G}_0$ is a singleton set. In this case, what do Theorems 5.3 and 5.5 tell us?

Theorem 5.6. *Let $\mathcal{G}_0$ be a singleton set $\{G_0\}$. The following are equivalent.*

(i) *$\mathcal{G}_0$ and $\mathcal{G}_1$ are (δ, c, r)-separable.*

(ii) *There is a sentence that is a conjunction of (δ, c, r)-sentences that separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*

(iii) *The sentence $\sigma(G_0, \delta, c, r)$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*

Proof. The fact that (i) and (iv) in the statement of Theorem 5.3 are equivalent implies immediately that (i) and (iii) are equivalent in the statement of Theorem 5.6. Clearly (iii) $\Rightarrow$ (ii) in the statement of Theorem 5.6, since $\sigma(G_0, \delta, c, r)$ is a conjunction of (δ, c, r)-sentences. Furthermore, (ii) $\Rightarrow$ (i) in the statement of Theorem 5.6, since (iii) $\Rightarrow$ (i) in the statement of Theorem 5.3. This concludes the proof. □

Theorem 5.7. *Let $\mathcal{G}_0$ be a singleton set $\{G_0\}$. The following are equivalent.*

(i) *$\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (δ, c, r)-separable.*

(ii) *There is a (δ, c, r)-sentence that separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*

(iii) *Some conjunct of $\sigma(G_0, \delta, c, r)$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$.*

Proof. This follows from Theorem 5.5 in the same way that Theorem 5.6 follows from Theorem 5.3. □

Intuitively, the fact that in the case of the Ajtai-Fagin (δ, c, r)-game (Theorem 5.7) a single (δ, c, r)-sentence separates $\mathcal{G}_0 = \{G_0\}$ from $\mathcal{G}_1$ reflects the fact that a single coloring of G_0 must simultaneously "work" against every member of $\mathcal{G}_1$. On the other hand, in the case of the original (δ, c, r)-game (Theorem 5.6), a conjunction of (δ, c, r)-sentences is required, because different colorings of G_0 might be needed by the spoiler for different choices of members of $\mathcal{G}_1$.

Note also another interesting difference between the separating sentences in the original (δ, c, r)-game (Theorems 5.3 and 5.6) versus those in the Ajtai-Fagin (δ, c, r)-game (Theorems 5.5 and 5.7). In the case of the original (δ, c, r)-game, the separating sentence depends only on $\mathcal{G}_0$ (that is, is independent of $\mathcal{G}_1$). By contrast, in the case of the Ajtai-Fagin (δ, c, r)-game, this sentence is not independent of $\mathcal{G}_1$. This is because in the Ajtai-Fagin game, a single coloring of G_0 must simultaneously work against every member of $\mathcal{G}_1$, and which coloring this is depends on $\mathcal{G}_1$.

The next proposition follows immediately from the definitions. It says, intuitively, that if the duplicator has a winning strategy in the original δ-ary NP game, then he has a winning strategy in the Ajtai-Fagin δ-ary NP game, with the same choices of graphs. This is what we would expect, since intuitively, it is even easier for the duplicator to win the Ajtai-Fagin game than the original game.

Proposition 5.8. *Let $\mathcal{G}_0$, $\mathcal{G}_1$ be classes of graphs. If $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-inseparable, then $\mathcal{G}_0$, $\mathcal{G}_1$ are Ajtai-Fagin (δ, c, r)-inseparable.*

As we shall see (Theorem 8.1), the converse is false. We are interested in comparing inseparability and Ajtai-Fagin inseparability to compare the graphs that can be used in a proof that a property is not in δ-ary NP using Ajtai-Fagin δ-ary NP games and using the original δ-ary NP games. An example of this reasoning appears in Section 7.

6 Ajtai-Fagin games are no stronger

The main theorem of this section (Theorem 6.2) is a strengthening of Theorem 5.4. It is partial converse to Proposition 5.8. It tells us that for each δ, c, r, there are c', r' such that if $\mathcal{G}_0$, $\mathcal{G}_1$ are Ajtai-Fagin (δ, c', r')-inseparable, then $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-inseparable. In fact, we can let $r' = r$. As we shall see (Theorem 8.1), we cannot always let $c' = c$. Hence, the converse of Proposition 5.8 is false, so we must settle for a partial converse.

The next proposition is a useful tool in proving Theorem 6.2.

Proposition 6.1. *The conjunction of a (δ, c_1, r)-sentence and a (δ, c_2, r)-sentence is equivalent to a single (δ, c_1c_2, r)-sentence.*

Proof. We first consider the case where c_1 and c_2 are each powers of 2 (say $c_1 = 2^{k_1}$ and $c_2 = 2^{k_2}$). In this case, a (δ, c_1, r) sentence is equivalent to a δ-ary Σ_1^1 sentence $\exists A_1 ... \exists A_{k_1} \psi_1$, where ψ_1 is an r-sentence. Similarly, a (δ, c_2, r) sentence is then equivalent to a δ-ary Σ_1^1 sentence $\exists A'_1 ... \exists A'_{k_1} \psi_2$, where ψ_2 is an r-sentence. We can assume without loss of generality that no A_i and A'_j are the same. Then the conjunction of the two sentences is equivalent to the sentence

$$\exists A_1 ... \exists A_{k_1} \exists A'_1 ... \exists A'_{k_2} (\psi_1 \wedge \psi_2),$$

which in turn is equivalent to a (δ, c_1c_2, r)-sentence.

In the general case (where c_1 and c_2 are not necessarily powers of 2), we use for the conjunction c_1c_2 colors, each of which is thought of as a pair (i, j) of colors, where i is one of the c_1 colors in the first sentence, and j is one of the c_2 colors in the second sentence. The details, which are simple but notationally tedious, are left to the reader. □

We now give the main theorem of this section.

Theorem 6.2. *Let δ, c, r be positive integers. There is c' such that for every $\mathcal{G}_0$, $\mathcal{G}_1$ that are Ajtai-Fagin (δ, c', r)-inseparable, also $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-inseparable.*

Proof. It follows from the proof of Theorem 4.1 that up to equivalence, there are only a finite number k of possible r-types (that involve only the binary relation symbol that represents the graph, and the "color" δ-ary relation symbols $B_1, \ldots, B_c$). This number k depends only on δ, c and r (in Section 9, we shall compute an upper bound on k, during the proof of Theorem 9.1). Therefore, up to equivalence, there are only k distinct (δ, c, r)-sentences of the form $\exists \{B_1, \ldots, B_c\} \psi$, where ψ is an r-type. Let $c' = c^k$. Assume that $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-separable; we shall show that $\mathcal{G}_0$, $\mathcal{G}_1$ are Ajtai-Fagin (δ, c', r)-separable.

From the fact that (i) $\Rightarrow$ (iv) in Theorem 5.3, the sentence $\bigvee_{G_0 \in \mathcal{G}_0} \sigma(G_0, \delta, c, r)$ separates $\mathcal{G}_0$ from $\mathcal{G}_1$. By Theorem 4.2, each sentence $\sigma(G_0, \delta, c, r)$ is of the form $\bigwedge_{i \in I} \exists\{B_1, \ldots, B_c\}\, \psi_i$, where each ψ_i is an r-type. Since, up to equivalence, there are at most k distinct r-types, it follows that each sentence $\sigma(G_0, \delta, c, r)$ is equivalent to the conjunction of at most k distinct (δ, c, r)-sentences. Therefore, by Proposition 6.1 applied repeatedly, $\sigma(G_0, \delta, c, r)$ is equivalent to a (δ, c', r)-sentence. Hence, a disjunction of (δ, c', r)-sentences separates $\mathcal{G}_0$ from $\mathcal{G}_1$. So from the fact that (ii) $\Rightarrow$ (i) in Theorem 5.5, where c' plays the role of c, it follows that $\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (δ, c', r)-separable, as desired. □

We can now prove Theorem 5.4 from the previous section.

Proof of Theorem 5.4. It follows immediately from Proposition 5.8 that (i) $\Rightarrow$ (ii). We now show that (ii) $\Rightarrow$ (i). Assume that (ii) holds, and let c, r be given. We wish to show that the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. Let c' be as in Theorem 6.2. By assumption, the duplicator has a winning strategy in the Ajtai-Fagin (δ, c', r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. That is, $\mathcal{G}_0$, $\mathcal{G}_1$ are Ajtai-Fagin (δ, c', r)-inseparable. So by Theorem 6.2, it follows that $\mathcal{G}_0$, $\mathcal{G}_1$ are (δ, c, r)-inseparable. That is, the duplicator has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$, as desired. □

Theorem 6.2 is rather powerful, since it guarantees the existence of a winning strategy for the duplicator in the original game (with a certain choice of parameters) given only the existence of a winning strategy for the duplicator in the Ajtai-Fagin game (with another choice of parameters). Intuitively, for a given choice of c, r, we use bigger graphs in the original (δ, c, r)-game than in the Ajtai-Fagin (δ, c, r)-game, since in the original game we use $\mathcal{G}_0$, $\mathcal{G}_1$ that correspond to the Ajtai-Fagin game with more colors. Note that the choice of c' is uniform, over all possible choices of $\mathcal{G}_0$, $\mathcal{G}_1$. As we shall see by example in the next section, Theorem 6.2 tells us that the same families of graphs can be used in the original game as in the Ajtai-Fagin game (such as to prove that a class is not in δ-ary NP). In the case of this example, we show how to extract from the proof of Theorem 6.2 a winning strategy for the duplicator in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$, including a coloring strategy.

7 Directed reachability

As we noted earlier, AJTAI and the author introduced their variation of monadic NP games in order to prove that directed reachability is not in monadic NP. In this section, we discuss this approach, and in particular discuss various senses in which the original monadic NP game is adequate and is not adequate to obtain this result.

Let $\delta = 1$, and let c and r be given. AJTAI and the author constructed (by probabilistic methods) a finite directed graph G_0 with points s, t where there is a directed path from s to t in G_0. In fact, G_0 consists of a path from s to t (these edges in the path are called "forward edges"), along with certain backedges. Thus, G_0 is (s, t)-connected. Denote the graph that is obtained by deleting the edge e from G_0 by $G_0 - e$. In particular, if e is a forward edge, then $G_0 - e$ is not (s, t)-connected. AJTAI and the author showed that however the spoiler colors the points of G_0 with the c colors, there is a forward edge e of G_0 such that when $G_1 = G_0 - e$ is colored in precisely the same way, point for point, as G_0 was colored, then the duplicator has a winning strategy in the r-game played on G_0 and G_1 (where, as before, the isomorphism must also respect color). Since G_0 is (s,t)-connected and G_1 is not, it follows from Theorem 3.4 that directed reachability is not in monadic NP.[7)]

Note that the duplicator does not commit himself to a choice of G_1 until the spoiler has committed himself to a coloring of G_0. This is the power of Ajtai-Fagin monadic NP games.

It is interesting to see what would happen if we were to try to use these pairs G_0, G_1 in the original monadic NP game rather than the Ajtai-Fagin monadic NP game. Intuitively, in the original game, the spoiler knows what G_0 and G_1 are before he colors G_0. This would be disastrous for a duplicator whose coloring strategy is to color G_1 by simply duplicating the coloring for G_0: if the spoiler knew which edge e were deleted from G_0 to form $G_1 = G - e$, this might dramatically influence his coloring of G_0 (for example, the spoiler might color the endpoints of e with special colors). In the Ajtai-Fagin monadic NP game, the spoiler must commit himself to a coloring of G_0 before he knows which edge e is deleted. This makes it easier for the duplicator to win.

In a proof that a problem is not in monadic NP using the original monadic NP game, we must give a coloring strategy for the duplicator, that tells the duplicator how to color G_1 as a function of the spoiler's coloring of G_0. This may be rather complicated. For example, in the author's original proof [6] that connectivity is not in monadic NP, it was shown, as we noted earlier, 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. Perhaps the hardest part of the proof lies in describing the duplicator's coloring strategy. By contrast, in AJTAI and the author's proof (using Ajtai-Fagin games) that directed reachability is not in monadic NP, describing the duplicator's coloring strategy is an easy step; the duplicator simply copies, point for point, the coloring of G_0. By making easier the task of finding a coloring strategy for the duplicator, we simplify our task of proving that a problem is not in monadic NP.

AJTAI and the author commented that they do not know how to prove their main result (that directed reachability is not in monadic NP) by using the original game. In such a proof, it would be necessary, given c, r, to show the existence of a pair $\widehat{G_0}$, $\widehat{G_1}$ of finite directed graphs where $\widehat{G_0}$ is (s,t)-connected, $\widehat{G_1}$ is not (s,t)-connected, and the duplicator has a winning strategy in the (c,r)-game over $\widehat{G_0}$, $\widehat{G_1}$. Since directed reachability is not in monadic NP (as AJTAI and the author showed), it follows from Theorem 3.2 that for each pair c, r, there is such a pair $\widehat{G_0}$, $\widehat{G_1}$. AJTAI and the author instead used Ajtai-Fagin monadic NP games, and worked with pairs G_0, G_1 where G_0 consists of a path from s to t, along with certain backedges, and where G_1 is the result of deleting some forward edge from G_0. AJTAI and the author said that it is not clear that such a pair G_0, G_1 could serve as $\widehat{G_0}$, $\widehat{G_1}$. We now show that it follows from Theorem 6.2 that indeed, such a pair G_0, G_1 could serve as $\widehat{G_0}$, $\widehat{G_1}$. This resolves AJTAI and the author's question.

[7)] ARORA and the author [2] show how to simplify AJTAI and the author's proof of this result. Both papers (AJTAI and FAGIN [1], and ARORA and FAGIN [2]) use exactly the same graphs G_0, G_1.

Theorem 7.1. *For every c, r, there is a graph G_0 that consists of a path from s to t, along with certain backedges, and a graph G_1 that is the result of deleting some forward edge from G_0, such that the duplicator has a winning strategy in the (c, r)-game over G_0, G_1.*

Proof. Assume that c, r are given. Let $\delta = 1$, and find c' as in Theorem 6.2. We know from [1] that there is G_0 that is a path from s to t, along with certain backedges, such that if $\mathcal{G}_0 = \{G_0\}$ and $\mathcal{G}_1$ consists of all graphs obtained from G_0 by deleting a single forward edge, then $\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (c', r)-inseparable. It follows from Theorem 6.2 that $\mathcal{G}_0$ and $\mathcal{G}_1$ are (c, r)-inseparable. Therefore, there is $G_1 \in \mathcal{G}_1$ such that the duplicator has a winning strategy in the (c, r)-game over G_0, G_1. □

It is important to note that we do *not* know a direct proof of Theorem 7.1, using only the original monadic NP game. There would be two phases in such a proof: the description of a strategy for the duplicator, and a proof that it is a winning strategy. In the description phase, we need in particular to describe how to select G_0, how to select G_1 (which differs from G_0 by having some forward edge deleted) and how the duplicator should color G_1 as a function of the spoiler's coloring of G_0.

It is perhaps instructive to see what our results tell us about how to select G_0, how to select G_1, and how the duplicator should color G_1 as a function of the spoiler's coloring of G_0. Let $\mathcal{G}_0$, $\mathcal{G}_1$ be as in the proof of Theorem 7.1. The choice of G_0 in the original (c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$ is of course determined as the unique member of the singleton set $\mathcal{G}_0$. We now discuss how to extract from the proof of Theorem 6.2 the way the duplicator should select G_1, along with a winning coloring strategy for the duplicator in the (c, r)-game over G_0, G_1, in the case of directed reachability.

Define a *c-coloring function* of a graph G to be a function associating with each of the nodes of G one of the c colors (for definiteness, we can think of the colors as being the numbers from 1 to c). Let C_1 and C_2 be c-coloring functions of G_0. Let H_1 (respectively H_2) be the colored graph that is the result of coloring G_0 with C_1 (respectively C_2). Let us say that the c-coloring functions C_1 and C_2 of G_0 are *r-equivalent over* G_0 if $H_1 \sim_r H_2$. Intuitively, two c-coloring functions are r-equivalent over G_0 if the duplicator has a winning strategy in the r-game over the colored graphs that are the result of coloring G_0 with the two c-coloring functions. It follows from Theorem 4.1 that there are only a finite number of r-equivalence classes of c-coloring functions of G_0, and that there is a finite upper bound on this number that is independent of the choice of G_0. Let $\mathcal{R}$ be a finite set of representatives of these equivalence classes. Thus, every c-coloring function of G_0 is r-equivalent over G_0 to some member of $\mathcal{R}$. The idea behind the proof of Theorem 6.2 is as follows. Let G_1 be the graph that the duplicator would select in the Ajtai-Fagin (c', r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$ if the spoiler were to color G_0 with the "hardest c'-coloring function of G_0", that is, a c'-coloring function of G_0 that encodes the simultaneous coloring of G_0 with each of the c-coloring functions of G_0 in $\mathcal{R}$. In the case (which we are focusing on) of showing that directed reachability is not in monadic NP, this graph G_1 is the result of deleting some forward edge from G_0, and in particular has the same set of nodes as G_0. We can now describe a winning coloring strategy for the duplicator in the (c, r)-game over G_0, G_1. Assume that the spoiler colors G_0 with the c-coloring function C_0. Let C_1 be a member of $\mathcal{R}$ that is r-equivalent over G_0 to C_0. Then the duplicator colors

G_1 with the c-coloring function C_1. Note that this coloring strategy is much more complicated than simply duplicating the coloring of G_0 on G_1 (the coloring strategy of the duplicator in the Ajtai-Fagin monadic NP game).

Remark 7.2. We seem to need to use Theorem 6.2, rather than the slightly weaker Theorem 5.4, to prove Theorem 7.1. What could we obtain by using only Theorem 5.4? Let $\mathcal{G}_0$ consist of all graphs that are a path from s to t, along with some backedges, and let $\mathcal{G}_1$ be the collection of all graphs that are the result of deleting an arbitrary forward edge from an arbitrary member of $\mathcal{G}_0$. We know from [1] that for every c', r', the duplicator has a winning strategy in the Ajtai-Fagin (c', r')-game over $\mathcal{G}_0, \mathcal{G}_1$. It follows from Theorem 5.4 that for every c, r, the duplicator has a winning strategy in the original (c, r)-game over $\mathcal{G}_0, \mathcal{G}_1$. Therefore, for every c, r, there are $G_0 \in \mathcal{G}_0$ and $G_1 \in \mathcal{G}_1$ such that the duplicator has a winning strategy in the (c, r)-game over G_0, G_1. But this result is not as strong as Theorem 7.1, since G_1 might be the result of deleting some forward edge from some member of $\mathcal{G}_0$ other than G_0.

8 Ajtai-Fagin games are stronger

In this section, we show that inseparability and Ajtai-Fagin inseparability are different, even in the monadic case $\delta = 1$. Thus, the converse of Proposition 5.8 is false. This tells us that there are situations where the spoiler requires strictly more colors to win the Ajtai-Fagin game than the original game. We discuss this more at the end of the section.

Theorem 8.1. *There are classes $\mathcal{G}_0$ and $\mathcal{G}_1$ of graphs, and constants c, r, such that $\mathcal{G}_0$ and $\mathcal{G}_1$ are (c, r)-separable, but $\mathcal{G}_0$ and $\mathcal{G}_1$ are Ajtai-Fagin (c, r)-inseparable.*

Proof. Let C_n be an undirected cycle with exactly n nodes. For definiteness, assume that the nodes are $\{1, \ldots, n\}$, and that there is an undirected edge between i and $i+1$ if $1 \leq i \leq n-1$, there is an undirected edge between n and 1, and there are no other edges. We refer to nodes with an edge between them as *neighbors*, and say that they are *adjacent*. Thus, each node has exactly two neighbors. Let $\mathcal{G}_0 = \{C_6\}$, and let $\mathcal{G}_1 = \{C_4, C_9\}$. We shall show that $\mathcal{G}_0, \mathcal{G}_1$ are $(2, 2)$-separable but Ajtai-Fagin $(2, 2)$-inseparable.

When we say that $\mathcal{G}_0, \mathcal{G}_1$ are $(2, 2)$-separable, we mean that if G_0 is C_6, and G_1 is either C_4 or C_9, then the *spoiler* has a winning strategy in the $(2, 2)$-game over G_0, G_1. Intuitively, if the spoiler knows what G_1 is (that is, either the spoiler knows that G_1 is C_4, or the spoiler knows that G_1 is C_9), then the spoiler can color G_0 in such a way as to guarantee a win.

By contrast, in the Ajtai-Fagin game, the duplicator has the option of seeing how the spoiler colors C_6 before deciding whether to continue the game by selecting C_4, or to continue the game by selecting C_9. In this case, as we shall show, it is the *duplicator* who has a winning strategy. That is, $\mathcal{G}_0, \mathcal{G}_1$ are Ajtai-Fagin $(2, 2)$-inseparable.

We first show that $\mathcal{G}_0, \mathcal{G}_1$ are $(2, 2)$-separable. To show this, we need only show that (1) the spoiler has a winning strategy in the $(2, 2)$-game over C_6, C_4, and (2) the spoiler has a winning strategy in the $(2, 2)$-game over C_6, C_9.

Let us assume for definiteness that the two colors are red and blue. To prove (2), we need only show that if G_0 is an even undirected cycle, and G_1 is an odd undirected cycle, then the spoiler has a winning strategy in the $(2,2)$-game over G_0, G_1. The spoiler colors G_0 by coloring even nodes red and odd nodes blue. Since G_1 has an odd number of nodes, it is easy to see that however the duplicator colors G_1, there will be two adjacent nodes of G_1 with the same color. In the 2-game that follows the coloring, the spoiler selects these two nodes of G_1 in his two moves. Since there are no two adjacent nodes in G_0 with the same color, it is clear that however the duplicator moves, the spoiler will win.

To prove (1), we need only show that if G_0 is an arbitrary undirected cycle (such as C_6) where the number of nodes is a multiple of 3, and if G_1 is an arbitrary undirected cycle where the number of nodes is *not* a multiple of 3 (such as C_4), then the spoiler has a winning strategy in the $(2,2)$-game over G_0, G_1. The spoiler colors G_0 by coloring the nodes red, red, blue, red, red, blue, red, red, blue, ..., with some integral number of repetitions of the pattern red, red, blue. We now show that since the number of nodes in G_1 is not a multiple of 3, however the duplicator colors G_1 the spoiler can win the remaining 2-game.

To show this, we first show that for the duplicator to have a chance to win, each node that is colored red in G_1 must have exactly one red neighbor. If the duplicator does not color G_1 so that each node that is colored red in G_1 has exactly one red neighbor, then there are two cases.

Case 1. Some node b_1 that is colored red in G_1 has no red neighbor. Then in the spoiler's first move in the remaining 2-game, he selects b_1 in G_1. The duplicator must then select a red node a_1 in G_0 to have a chance to win. The spoiler then selects a red neighbor of a_1 in G_0 (this is guaranteed to exist by the spoiler's coloring of G_0). The duplicator cannot then select a red neighbor of b_1 in G_1, so the duplicator loses.

Case 2. Some node b_1 that is colored red in G_1 has two red neighbors. Then in the spoiler's first move in the remaining 2-game, the spoiler selects b_1 in G_1. The duplicator must then select a red node a_1 in G_0 to have a chance to win. The spoiler then selects a blue neighbor of a_1 in G_0 (this is guaranteed to exist by the spoiler's coloring of G_0). The duplicator cannot then select a blue neighbor of b_1 in G_1, so the duplicator loses.

We just showed that for the duplicator to have a chance to win, each node that is colored red in G_1 must have exactly one red neighbor. By a similar argument, each node that is colored blue in G_1 must have *every* neighbor be red. It is easy to see from these two facts that for the duplicator to have a chance to win, he must color G_1 starting at some node and going around the cycle with the colors red, red, blue, red, red, blue, red, red, blue, ..., with some integral number of repetitions of the pattern red, red, blue. But this is impossible if the number of nodes in G_1 is not a multiple of 3.

We just showed that $\mathcal{G}_0$, $\mathcal{G}_1$ are $(2,2)$-separable. We now show that $\mathcal{G}_0$, $\mathcal{G}_1$ are Ajtai-Fagin $(2,2)$-inseparable. We consider every possible way that the spoiler can color C_6. There are three cases.

Case 1. The spoiler colors C_6 with the two colors alternating red, blue, red, blue, red, blue. Then the duplicator selects C_4 from $\mathcal{G}_1$, and colors C_4 with the two colors alternating red, blue, red, blue. It is straightforward to verify that the duplicator can now win the remaining 2-game.

Case 2. The spoiler colors C_6 so that there are at least three consecutive nodes with the same color (three distinct nodes a, b, c are *consecutive* if b has a and c as its neighbors). Say for definiteness that nodes 1,2,3 are each colored red. Then the duplicator selects C_9 from $\mathcal{G}_1$, and colors C_9 by coloring the nodes 1,2,3,4,5,6 in C_9 each red, and coloring the nodes 7,8,9 in C_9 the same colors as the nodes 4,5,6, respectively, are colored in C_6. We leave to the reader the fairly easy verification that the duplicator then has a winning strategy in the remaining 2-game. The key point is that if the spoiler picks on his first move of the remaining 2-game any of the nodes 2,3,4,5 in C_9, then the duplicator picks node 2 in C_6 on his first move; if the spoiler picks on his first move of the remaining 2-game any of the nodes 1,6,7,8,9, respectively, in C_9 on his first move, then the duplicator picks 1,3,4,5,6, respectively, in C_6 on his first move of the 2-game.

Case 3. The spoiler colors C_6 so that the maximal number of consecutive nodes with the same color is exactly two. Let us say for definiteness that nodes 1 and 2 are colored blue in C_6. Then nodes 3 and 6 must be colored red in C_6, or else there would be three consecutive blue nodes in C_6. There are now four subcases, depending on how nodes 4 and 5 are colored in C_6.

Subcase 3a. Nodes 4 and 5 are both colored red in C_6. This is impossible, since nodes 3,4,5,6 would all be colored red, which contradicts the assumption that the maximal number of consecutive nodes with the same color is two.

Subcase 3b. Node 4 is colored red in C_6, and node 5 is colored blue in C_6. Then the duplicator selects C_9 from $\mathcal{G}_1$, and colors C_9 by coloring the nodes $1, \ldots, 9$ as blue, blue, red, red, blue, blue, red, blue, red. It is not hard to verify that the duplicator has a winning strategy in the remaining 2-game. For example, if on his first move in the remaining 2-game the spoiler selects a red node in one of the graphs, whose neighbors are both blue (respectively, red and blue), then the duplicator can do the same in the other graph, and then the duplicator can win no matter what the spoiler's next move is. Similarly, if on his first move in the remaining 2-game the spoiler selects a blue node in one of the graphs, whose neighbors are both red (respectively, red and blue), then the duplicator can do the same in the other graph, and then the duplicator can win no matter what the spoiler's next move is.

Subcase 3c. Node 4 is colored blue in C_6, and node 5 is colored red in C_6. After we reverse the colors red and blue, the coloring of C_6 has the same pattern as in as Subcase 3b.

Subcase 3d. Nodes 4 and 5 are both colored blue in C_6. Thus, C_6 is colored blue, blue, red, blue, blue, red (that is, with the pattern blue, blue, red repeated twice). Then the duplicator selects C_9 from $\mathcal{G}_1$, and colors C_9 blue, blue, red, blue, blue, red, blue, blue, red (that is, with the pattern blue, blue, red repeated three times). It is easy to see that the duplicator has a winning strategy in the remaining 2-game.

This concludes the demonstration that $\mathcal{G}_0, \mathcal{G}_1$ are $(2,2)$-separable, but are Ajtai-Fagin $(2,2)$-inseparable. □

The contrapositive of Theorem 6.2 tells us that since $\mathcal{G}_0, \mathcal{G}_1$ in the proof of Theorem 8.1 are $(2,2)$-separable, there is c' such that $\mathcal{G}_0, \mathcal{G}_1$ are Ajtai-Fagin $(c',2)$-separable. The proof of Theorem 8.1 shows that $c' > 2$. It is not hard to verify that in this case, we can take $c' = 3$ (the spoiler colors C_6 red, red, blue, blue, green, green). So in this example, 3 colors are required for the spoiler to have a winning strategy in the Ajtai-Fagin game, but only 2 colors to have a winning strategy in the original game. Hence, in this case, the Ajtai-Fagin game is harder for the spoiler to win (and therefore easier for the duplicator to win) than the original game.

9 How many more colors are required?

The contrapositive of Theorem 6.2 says that for each choice of the arity δ, the number c of colors and the number r of rounds, there is c' such that whenever $\mathcal{G}_0, \mathcal{G}_1$ are (δ, c, r)-separable, then $\mathcal{G}_0, \mathcal{G}_1$ are Ajtai-Fagin (δ, c', r)-separable. Let us denote the minimal such value of c' by $F(\delta, c, r)$. Intuitively, when the spoiler can win the original δ-ary NP game with c colors and r rounds, then $c' = F(\delta, c, r)$ is the number of colors the spoiler needs to win the Ajtai-Fagin game. We saw in Theorem 8.1 that there are δ, c, r where $F(\delta, c, r) > c$. This says that the spoiler requires strictly more colors to win the Ajtai-Fagin game than the original game. How many more colors are required? That is, how much bigger does $F(\delta, c, r)$ need to be than c? We now give an upper bound on $F(\delta, c, r)$. Define f by letting

$$f(0) = 2r^2 + cr^\delta \quad \text{and} \quad f(m+1) = 2^{f(m)} \text{ for each } m.$$

Theorem 9.1. $F(\delta, c, r) \leq c^{f(r+1)}$.

Proof. As we see from the proof of Theorem 4.1, there are $2r^2 + cr^\delta$ atomic formulas, with r^2 of the form $v_i = v_j$, with r^2 of the form Pv_iv_j, and with cr^δ of the form $B_\ell v_{i_1} \ldots v_{i_\delta}$. It follows easily from the proof of Theorem 4.1 that for each m with $0 \leq m \leq r$, the number of m-types is at most $f(m+1)$. In particular, the number of r-types is at most $f(r+1)$. It then follows from the proof of Theorem 6.2 that we can take c' to be $c^{f(r+1)}$. Therefore, $F(\delta, c, r) \leq c^{f(r+1)}$. □

Note that the upper bound $c^{f(r+1)}$ in Theorem 9.1 contains a tower of $r+2$ exponents, where the top exponent is a polynomial in r, c, and r^δ. This represents a nonelementary growth rate.

10 Conclusions and open problems

The original δ-ary NP game and the Ajtai-Fagin δ-ary NP game are known to be equivalent in a "global" sense. Thus, given δ, the duplicator has a winning strategy in the original game for each choice of the remaining parameters (number of colors and number of rounds) if and only if the duplicator has a winning strategy in the Ajtai-Fagin game for each choice of the remaining parameters. In this paper, we investigated the "local" aspects of this equivalence. First, we considered the families of graphs used in the games. We showed each family of graphs used in the Ajtai-Fagin game to prove that a problem is not in δ-ary NP can in principle be used in the original game to prove the same result (where for a given choice of parameters, bigger graphs of the same type are used for the original game than for the Ajtai-Fagin game). To obtain this result, we obtained strengthened versions of the equivalence between the original game and the Ajtai-Fagin game. Second, we considered the number of colors required for the spoiler to win when the choices of graphs are fixed.

We showed that there are situations where the spoiler requires strictly more colors to win the Ajtai-Fagin game than the original game. In this sense, the Ajtai-Fagin game is strictly stronger than the original game. Our analysis gives a nonelementary upper bound on the number of extra colors that are required for the spoiler to win the Ajtai-Fagin game than the original game.

There are a number of open problems concerning the behavior of $F(\delta, c, r)$, as defined in Section 9. It follows from Theorem 8.1 that there are δ, c, r such that $F(\delta, c, r) > c$, which corresponds to the fact that there are $\mathcal{G}_0$, $\mathcal{G}_1$, r such that the spoiler requires strictly more colors to win the Ajtai-Fagin game over $\mathcal{G}_0$, $\mathcal{G}_1$ than the original game. That is, the duplicator has a winning strategy in the Ajtai-Fagin (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$, but the spoiler has a winning strategy in the original (δ, c, r)-game over $\mathcal{G}_0$, $\mathcal{G}_1$. In fact, we showed that $F(1,2,2) > 2$ (we note that it is easy to see that $F(1, c, r) = c$ if $c = 1$ or $r = 1$, so this is a minimal example). We now list some open problems.

· Is $F(\delta, c, r) > c$ for every δ, c, r with $\delta > 1$, and with $c > 2$ and $r > 2$? (We conjecture that the answer is "Yes".)

· What is the growth rate of $F(\delta, c, r)$? In particular, is there a nonelementary lower bound to go along with the nonelementary upper bound given in Theorem 9.1? (Again, we conjecture that the answer is "Yes".)

· How do the answers to these questions change if we restrict our attention to pairs $\mathcal{G}_0$, $\mathcal{G}_1$ where $\mathcal{G}_0 = \mathcal{C}$ and $\mathcal{G}_1 = \overline{\mathcal{C}}$ for some class $\mathcal{C}$? (We conjecture that the answers do not change.)

References

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

[2] ARORA, S., and R. FAGIN, On winning strategies in Ehrenfeucht-Fraïssé games. Research Report RJ 9833, IBM, June 1994; to appear in Theoretical Computer Science.

[3] DURAND, A., C. LAUTEMANN, and T. SCHWENTICK, Subclasses of binary NP. To appear in Journal of Logic and Computation.

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

[5] FAGIN, R., Generalized first-order spectra and polynomial-time recognizable sets. In: Complexity of Computation (R. M. KARP, ed.), SIAM-AMS Proceedings **7** (1974), 43 - 73.

[6] FAGIN, R., Monadic generalized spectra. Zeitschrift Math. Logik Grundlagen Math. **21** (1975), 89 - 96.

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

[8] FAGIN, R., L. STOCKMEYER, and M. Y. VARDI, On monadic NP versus monadic co-NP. Information and Computation **120** (1995), 78 - 92.

[9] IMMERMAN N., Descriptive and computational complexity. In: Computational Complexity Theory (J. HARTMANIS, ed.), Proc. AMS Symp. in Appl. Math. **38** (1989), 75 - 91.

(Received: July 15, 1996)