THE JOURNAL OF SYMBOLIC LOGIC
Volume 55, Number 1, March 1990

REACHABILITY IS HARDER FOR DIRECTED THAN FOR UNDIRECTED FINITE GRAPHS

MIKLOS AJTAI AND RONALD FAGIN

Abstract. Although it is known that reachability in undirected finite graphs can be expressed by an existential monadic second-order sentence, our main result is that this is not the case for directed finite graphs (even in the presence of certain "built-in" relations, such as the successor relation). The proof makes use of Ehrenfeucht-Fraïssé games, along with probabilistic arguments. However, we show that for directed finite graphs with degree at most k, reachability *is* expressible by an existential monadic second-order sentence.

§1. Introduction. If s and t denote distinguished points in a directed (resp. undirected) graph, then we say that a graph is *(s, t)-connected* if there is a directed (undirected) path from s to t. We sometimes refer to the problem of deciding whether a given directed (undirected) graph with two given points s and t is (s, t)-connected as the directed (undirected) *reachability* problem.

Consider the undirected graphs in Figures 1 and 2. It is easy to tell at a glance that the graph in Figure 1 is (s, t)-connected (since s and t are in the same connected component), and that the graph in Figure 2 is not (s, t)-connected (since s and t are in different connected components). Consider now the directed graphs in Figures 3 and 4. As the reader can verify, the graph in Figure 3 is (s, t)-connected, and the graph in Figure 4 is not. However, in the case of Figures 3 and 4, it is no longer possible to tell at a glance.

Of course, "tell at a glance" is hardly a precise technical notion. As we shall discuss in §2, researchers in computational complexity have struggled, so far unsuccessfully, to prove that on general-purpose models of computation (such as Turing machines), the reachability problem is harder for directed graphs than for undirected graphs. We prove that, in a certain precise sense, the directed case is indeed harder than the undirected case. The distinction we make is in terms of expressibility, as we now explain.

We begin with a few conventions. Since we are concerned only with *finite* graphs, whenever we say "graph", we mean "finite graph". Also, since a theme of this paper is

Received July 5, 1988; revised January 31, 1989.
A preliminary version of this paper appeared in the Proceedings of the 1988 IEEE Symposium on Foundations of Computer Science, pp. 358–367.

0022-4812/90/5501-0009/$04.80

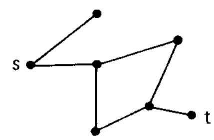

FIGURE 1

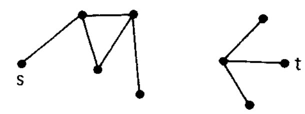

FIGURE 2

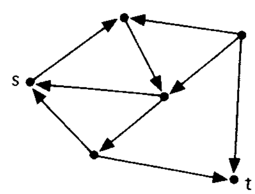

FIGURE 3

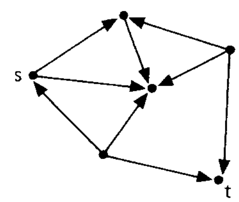

FIGURE 4

to contrast the situation for directed versus undirected graphs, whenever we say simply "graph", we mean either directed or undirected graph; when it is important to distinguish whether the graph is directed or undirected, we shall do so. For convenience, we consider only irreflexive graphs, that is, graphs where there is no edge from some point to itself. If G is a graph 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 $\exists A_1 \cdots \exists A_k \psi$, 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 P) is 3-colorable. In this sentence, the three colors are represented by A_1, A_2, and A_3. Let ψ_1 say "Each point has exactly one color". Thus, ψ_1 is

$$\begin{aligned}\forall x((A_1x \wedge \neg A_2x \wedge \neg A_3x) \vee (\neg A_1x \wedge A_2x \wedge \neg A_3x)\\ \vee (\neg A_1x \wedge \neg A_2x \wedge A_3x)).\end{aligned}$$

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

$$\begin{aligned}\forall x \forall y((A_1x \wedge A_1y \Rightarrow \neg Pxy) \wedge (A_2x \wedge A_2y \Rightarrow \neg Pxy)\\ \wedge (A_3x \wedge A_3y \Rightarrow \neg Pxy)).\end{aligned}$$

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

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

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. A class $\mathscr{C}$ of graphs is said to be (monadic) Σ_1^1 if it is the class of all graphs that obey some fixed (monadic) Σ_1^1 sentence. A (monadic) Σ_1^1 class is also called a (monadic) generalized spectrum. As we have just seen, the class of 3-colorable graphs is monadic Σ_1^1.

As another example, we now show that the class of graphs that are not connected is monadic Σ_1^1 (this demonstration is from [Fa2]). Let ψ_1 say "The set A is nonempty and its complement is nonempty", that is,

$$\exists x \exists y (Ax \wedge \neg Ay).$$

Let ψ_2 say "There is no edge between A and its complement", that is,

$$\forall x \forall y (Ax \wedge \neg Ay \Rightarrow \neg Pxy).$$

It is clear that the monadic Σ_1^1 sentence $\exists A(\psi_1 \wedge \psi_2)$ characterizes the class of undirected graphs that are not connected. (A directed graph is said to be connected if the undirected version, where we ignore the directions on the edges, is connected; we can then show that the class of nonconnected *directed* graphs is monadic Σ_1^1 by replacing $\neg Pxy$ in ψ_2 by $\neg Pxy \vee \neg Pyx$.)

Let us define a class to be (monadic) Π_1^1 if its complement is (monadic) Σ_1^1. Fagin [Fa1] showed that, in a precise sense, Σ_1^1 is the same as NP (the class of languages recognizable nondeterministically in polynomial time [GJ]). In particular, the question as to whether $\Sigma_1^1 = \Pi_1^1$ (which would imply, for example, that 3-colorability is Π_1^1) is equivalent to the famous problem of whether NP = co-NP. Furthermore, it follows from the theory of NP-completeness [GJ] and from the equivalence of Σ_1^1 and NP that $\Sigma_1^1 = \Pi_1^1$ if and only if 3-colorability is Π_1^1 [Fa1]. (We can replace "3-colorability" in the previous statement by any other NP-complete problem on graphs, such as Hamiltonicity.) Although it is an open problem as to whether $\Sigma_1^1 = \Pi_1^1$, Fagin [Fa2] showed that *monadic* Σ_1^1 is different from *monadic* Π_1^1. In particular, he showed that *connectivity* (i.e., the class of connected graphs) is not monadic Σ_1^1, although it is monadic Π_1^1 (since, as we saw above, nonconnectivity is monadic Σ_1^1). This is true whether we consider directed or undirected graphs. This result was generalized by de Rougemont [Ro], by showing that it holds even if there is a built-in successor relation.

Since undirected connectivity is not monadic Σ_1^1, it came as somewhat of a surprise when Kanellakis ([Ka1]; see also [BKBR]) observed that undirected reachability is monadic Σ_1^1 (if we were to allow also infinite graphs, then this result would be false, as we can see by an easy compactness argument). To see that undirected reachability is monadic Σ_1^1, let ψ_1 be $As \wedge At$, that is, "The set A contains both s and t"; let ψ_2 be

$$\exists x \forall y ((Ay \wedge Psy) \Leftrightarrow x = y),$$

that is, "s has an edge to exactly one member of A"; let ψ_3 be

$$\exists x \forall y ((Ay \wedge Pty) \Leftrightarrow x = y),$$

that is, "t has an edge to exactly one member of A"; and let ψ_4 be

$$\begin{aligned}&\forall x((Ax \wedge (x \neq s) \wedge (x \neq t))\\ &\qquad \Rightarrow (\exists y_1 \exists y_2 (y_1 \neq y_2 \wedge \forall z ((Az \wedge Pxz) \Leftrightarrow (z = y_1 \vee z = y_2))))),\end{aligned}$$

that is, "Every member of A except for s and t has an edge to precisely two members of A". If ψ is taken to be $\psi_1 \wedge \psi_2 \wedge \psi_3 \wedge \psi_4$, then, as we now show, the Σ_1^1 sentence $\exists A\psi$ says that the graph is (s, t)-connected. For, if the graph is (s, t)-connected, and if A is taken to consist of those vertices on a shortest path from s to t, then ψ holds. Conversely, if ψ holds, then there is a path starting at s that passes through only vertices in A. The path must end somewhere, since the graph is finite; however, the only place it can end is at t. So the graph is (s, t)-connected.

We have just seen that undirected reachability is monadic Σ_1^1. Why can't a sentence in this spirit be used to show that *directed* reachability is Σ_1^1? The problem lies in "backedges". For example, consider a directed graph where the shortest (directed) path from s to t is on a path from s to a to b to c to t. Assume also that there is an edge (a "backedge") from c to a. Then the natural directed analogue of Kanellakis' Σ_1^1 sentence $\exists A\psi$ above fails, since if A were taken to be the set $\{s, a, b, c, t\}$, then in addition to the outgoing edge from a to b, there are two (not one) incoming edges to a from members of A.

Kanellakis posed as an open problem [Ka1] the question of whether directed reachability is monadic Σ_1^1. We show that it is not. However, interestingly enough, we show that for each positive integer k, the class of directed (s, t)-connected graphs where the indegree and outdegree of each vertex is at most k is monadic Σ_1^1.

In fact, we prove even stronger results. We show that our result that directed reachability is not monadic Σ_1^1 still holds, even in the presence of built-in relations from a large class, which includes the successor relation. In §3, we explain what it means for built-in relations to be present, and tell why it is of interest.

As we shall discuss in detail shortly, the nonexpressibility results of [Fa2] and [Ro] are obtained by considering Ehrenfeucht-Fraïssé-type games, and showing that one player ("the duplicator") has a winning strategy.[1] The graphs used in [Fa2] and [Ro] are explicitly described (in fact, the graphs in [Fa2] are just disjoint unions of cycles). Our approach has several novel features, two of which we now mention. First, the construction of the graphs we use is probabilistic, rather than deterministic. Second, we introduce a new game which, on the face of it, is easier for the duplicator to win, and prove that with high probability the duplicator does indeed have a winning strategy. This is sufficient to conclude our nonexpressibility results.

We consider our proof techniques to be of independent interest. It follows from Theorem 1 of [Fa2] that if Σ_1^1 is not closed under complement, then this can be proven by using an Ehrenfeucht-Fraïssé-type game argument. Hence, our new proof techniques may be a step on the road towards showing that Σ_1^1 is not closed under complement (or equivalently, by [Fa1], that NP is not closed under complement).

We now state explicitly our main theorem. For ease in description, we defer until later the statement as to how we can extend our main theorem by allowing certain built-in relations, such as the successor relation.

THEOREM 1.1. *Directed reachability is not monadic* Σ_1^1.

The structure of our proof of Theorem 1.1 is as follows. Let φ be a monadic Σ_1^1 sentence $\exists A_1 \cdots \exists A_k \psi(A_1, \ldots, A_k)$ that allegedly characterizes directed (s, t)-connectivity (where we have suppressed mention of s, t, and the graph predicate P in

[1] Following Joel Spencer, 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".

ψ for convenience). We construct (by probabilistic methods) a directed graph G with points s, t where there is a directed path from s to t in G. Thus, G is (s, t)-connected, so G satisfies φ. Hence, there are subsets $A_1, \ldots, A_k$ of the vertices of G such that G satisfies $\psi(A_1, \ldots, A_k)$. Let us denote the graph that is obtained by deleting the edge e from G by $G - e$. We show that there is an edge e of G such that (a) there is no directed path from s to t in $G - e$, and (b) $G - e$ satisfies $\psi(A_1, \ldots, A_k)$. This is a contradiction, since $G - e$ then satisfies φ, but is not (s, t)-connected.

In fact, since it is easy to see that reachability is monadic Π^1_1 in both the directed and undirected cases (by almost the same argument as we gave earlier that connectivity is monadic Π^1_1), it follows that undirected reachability is monadic $\Sigma^1_1 \cap \Pi^1_1$, while directed reachability is monadic Π^1_1 but not monadic Σ^1_1.

The next theorem says that for bounded degree graphs, directed reachability *is* monadic Σ^1_1. The proof of this theorem is in §5.

THEOREM 1.2. *Let k be a positive integer. The class of directed (s, t)-connected graphs where the indegree and outdegree of each vertex is at most k is monadic Σ^1_1.*

We note that the monadic Σ^1_1 sentence that we use to prove Theorem 1.2 has $O(k^2)$ existentially quantified monadic relation symbols.

We now describe the organization of the paper. In §2, we consider differences (and possible differences) in the computational complexity of problems on directed versus undirected graphs. In particular, we discuss the body of empirical evidence that in general-purpose models of computation, reachability is harder for directed graphs than for undirected graphs. In §3, we explain the meaning and significance of allowing built-in relations. In §4, we explain Ehrenfeucht-Fraïssé-type games that others have used, and tell how we modify the approach. We also give the result (Theorem 4.6) from which our nonexpressibility results are obtained. In §5, we prove that directed reachability is monadic Σ^1_1 if we restrict our attention to bounded-degree graphs. In §6, we show that a natural modification of the construction used to prove nonexpressibility does not work. In §7, we give a variation, that we shall utilize, of the well-known result that the probability that the number of successful independent trials differs from the mean by more than a constant times the mean is exponentially small. In §8, we give the proof of Theorem 4.6, which, as we noted, implies our nonexpressibility results.

§2. Computational complexity issues. In this section, we discuss differences (and possible differences) between the computational complexity of problems on directed versus undirected graphs. In particular, our main focus is reachability.

There are various cases where a problem is in some sense "harder" for directed graphs than for undirected graphs. For example, consider the kernel problem (problem GT57 in Garey and Johnson [GJ]). An *independent set* in a graph is a set of points with no edges between them. A *kernel* of a graph is an independent set S of vertices such that for every point y not in S there is a member x of S where there is an edge from x to y. Chvátal [Chv] shows that the kernel problem (the problem of deciding whether a given graph has a kernel) is NP-complete for directed graphs. However, for undirected graphs, the problem is trivial: the answer is always "yes" (since in an undirected graph, a kernel is just a maximal independent set).

There are other problems where the directed case is provably harder than the undirected case, including VLSI area bounds for doing on-line depth-first search (where there is an $\Omega(n^2)$ lower bound in the directed case [HMS], and an $O(n^{5/3})$ upper bound in the undirected case [An]). Further, there are problems where the undirected case seems to be (but has never been proven to be) harder, including parallel ear decomposition [Lo] and parallel depth-first search [AA].

We now turn to another example of the latter phenomenon, which is of greatest interest for this paper, namely reachability. Savitch [Sa] shows that directed reachability is nondeterministic log-space complete with respect to log-space reductions. However, it is unknown whether the same is true for undirected reachability. As another possible difference, Aleliunas et al. [AKLLR] show that there is an $O(\log n)$ space nonuniform algorithm for undirected reachability (with a polynomial in n amount of "advice" on the tape). However, the best that is known for directed reachability is $O(\log^2 n)$ space (this follows from Savitch's theorem [Sa]). Further, Aleliunas et al. show that undirected reachability can be solved by a probabilistic Turing machine (with small probability of error) in space $O(\log n)$ and polynomial expected time. This was recently improved by Borodin et al. [BCDRT] into an errorless probabilistic algorithm with the same space and time bounds. Again, it is not known whether there is such a probabilistic algorithm (even allowing a small probability of error) for directed reachability. Recently, Karchmer and Wigderson [KW] proved an $\Omega(\log^2 n/\log\log n)$ lower bound on the depth of a monotone circuit that tests reachability on n-vertex graphs. Their lower bound applies in both the directed and undirected cases.

There are also apparent (but again unproven) differences between directed and undirected reachability when we consider parallel computation. A major problem in parallel complexity (which is referred to by Ullman [Ul]) is that there is no known NC algorithm for directed reachability that uses fewer than roughly n^α processors, where α is the exponent for the time to multiply matrices (by [CW], we know that $2 \leq \alpha < 2.376$). The best NC algorithm (in terms of number of processors) that is known for directed reachability is by Pan and Reif [PR]; their algorithm uses essentially n^α processors and runs in $O(\log^2 n)$ time. By contrast, Cole and Vishkin [CV] show that undirected reachability can be solved in $\log n$ time with $n \log^* n/\log n$ processors. Resolving the parallel complexity of directed reachability is an important theoretical problem for parallel logic programming [BKBR], [Ka2].

Although it has never been proven that directed reachability is harder than undirected reachability in any general-purpose model of computation, there are certain proven distinctions between the directed and undirected cases. For example, starting at a given vertex in an undirected graph with n vertices, the expected time to visit all of the vertices is $O(n^3)$ [AKLLR]. In fact, this is the basis of the algorithm of [AKLLR] mentioned earlier for determining (s, t)-connectivity by a probabilistic Turing machine with small probability of error in space $O(\log n)$ and polynomial expected time: a random walk is started at s and continued for $O(n^3)$ steps, and we see whether or not t is ever reached. However, as noted in [AKLLR], there are *directed* graphs where the expected time to visit all of the vertices is exponential. Another distinction between directed and undirected reachability has been proven on special-purpose machines called *JAG*'s (*Jumping Automata for Graphs*). A JAG can move pebbles from a limited supply along the edges of a graph under finite-state control; the machine can detect when two pebbles coincide, and can cause one pebble to jump to another. Cook and Rackoff [CR] prove that directed reachability on a JAG requires space $\Omega(\log^2 n/\log\log n)$. This result was extended to *RJAG*'s (*randomized JAG*'s) by Berman and Simon [BS], if the RJAG is required to be both reliable (probability of acceptance bounded away from 1/2) and fast (polynomial time). However, the result of [AKKLR] on random walks mentioned above implies that such RJAG's can determine undirected reachability in space $O(\log n)$.

Our results show that in terms of expressibility, directed reachability is harder than undirected reachability: undirected reachability is monadic Σ^1_1, but directed reachability is not. In fact, as we noted earlier, undirected reachability is monadic $\Sigma^1_1 \cap \Pi^1_1$, while directed reachability is monadic Π^1_1 but not monadic Σ^1_1.

Unfortunately, our nonexpressibility result does not seem to translate into a lower bound on computational complexity. Thus, our results do not give us a proof that directed reachability is harder in some computational complexity sense than undirected reachability. There are two reasons for this, which we now discuss.

In terms of computational complexity, the "hard" graph problems in NP are the NP-complete problems. In terms of expressibility (at least as far as this paper is concerned), the "hard" graph problems in NP (i.e., the hard graph problems in Σ^1_1) are those that are not monadic Σ^1_1. However, these notions of "hardness" are orthogonal. Thus, there are NP-complete problems (such as 3-colorability) that are monadic Σ^1_1, whereas there are problems such as directed reachability that are easy in terms of computational complexity but that are not monadic Σ^1_1.

The second reason that our results do not translate into lower bounds on computational complexity is that, except for the issue of the number of processors required by an NC algorithm for directed reachability, the computational complexity of directed reachability in each of the senses we have discussed is reducible to the computational complexity of directed reachability for graphs of indegree and outdegree at most two. This is because if, for example, the vertex u has outdegree bigger than two, and if S is the set of vertices v such that $\langle u, v\rangle$ is an edge, then we can interpolate a tree with root u and with S as its set of leaves. So for reachability, many natural computational complexity measures do not distinguish between directed graphs of bounded and unbounded degree. However, as we show in this paper, bounded degree directed reachability is monadic Σ^1_1, whereas unbounded degree directed reachability is not.

§3. Allowing built-in relations. As we mentioned in the Introduction, we prove that directed reachability is not monadic Σ^1_1, even in the presence of "built-in" relations from a large class, which includes the successor relation. We now explain the meaning and significance of allowing built-in relations.

If $V = \{v_1, \ldots, v_n\}$ is a set of points, then a *successor relation* (over V) is a binary relation S with universe (set of vertices) V such that there is an ordering $v_1 < v_2 < \cdots < v_n$ of V where $S = \{\langle v_1, v_2\rangle, \langle v_2, v_3\rangle, \ldots, \langle v_{n-1}, v_n\rangle\}$. Other minor variations on this definition are sometimes used. The following theorem shows that our main result can be extended to allowing a built-in successor relation.

THEOREM 3.1. *Directed reachability is not monadic Σ^1_1, even in the presence of a built-in successor relation.*

A built-in successor relation is a successor relation that is "attached to" each universe V. Thus, having a built-in successor relation means that there is a fixed "successor assignment function" $\hat{S}$ that associates with each universe V a successor relation $\hat{S}(V)$ over V.[2] So Theorem 3.1 says that there is no monadic Σ_1^1 sentence $\varphi(P, s, t, S)$ and successor assignment function $\hat{S}$ such that a graph $G = \langle V, P, s, t \rangle$ (with universe V, binary relation P over V, representing the edges, and distinguished points s and t) is (s, t)-connected iff $\langle V, P, s, t, \hat{S}(V) \rangle \models \varphi$. For notational convenience, we shall not distinguish between, say, a binary relation symbol P and the binary relation which is the interpretation of P. It is important to notice that the successor relation $\hat{S}(V)$ depends only on the universe V, and not on the whole graph $G = \langle V, P, s, t \rangle$; otherwise, it is easy to see that Theorem 3.1 would be false (since the successor relation could "encode the information" in a trivial way as to whether or not the graph is (s, t)-connected).

There is another notion of "in the presence of a built-in successor relation". Using this notion, "directed reachability is monadic Σ_1^1 in the presence of a built-in successor relation" would mean that there is a monadic Σ_1^1 sentence $\varphi(P, s, t, S)$ such that if $\langle V, P, s, t, S \rangle$ is a graph as before, with a successor relation, then $\langle V, P, s, t \rangle$ is (s, t)-connected iff $\langle V, P, s, t, S \rangle \models \varphi$. It is not hard to verify that these notions of "in the presence of a built-in successor relation" are equivalent (the proof depends on the fact that all successor relations over V are isomorphic). We have chosen our formulation for reasons that will be apparent shortly.

We now comment on why it is of interest to allow a built-in successor relation. It is very natural in computer science to allow a built-in successor relation, since actual data as stored in a computer always has an implicit ordering (such as lexicographic ordering), based on the computer representation. Allowing a built-in successor relation can sometimes make a big difference. For example, Immerman [Im2] and Vardi [Va] independently show that a property is polynomial-time recognizable iff it can be expressed by a fixpoint sentence with successor. Allowing successor is crucial in this case, since Chandra and Harel [CH] show that without successor, there is no fixpoint sentence for evenness ("the number of points is even"). It follows from results of de Rougemont [Ro] and Immerman [Im1] that Hamiltonicity is not expressible by a fixpoint sentence. If this result could be extended to show that Hamiltonicity is not expressible by a fixpoint sentence *with successor*, then by Immerman and Vardi's characterization, this is equivalent to showing that $\mathrm{P} \neq \mathrm{NP}$!

As another example, Kolaitis and Vardi [KV] consider a class called "strict Σ_1^1", which are properties describable by Σ_1^1 sentences where the first-order part has a certain restricted quantifier structure. Kolaitis and Vardi prove that strict Σ_1^1 is not closed under complement. It follows easily from the construction in [Fa1] that strict Σ_1^1 with successor is, in a precise sense, equal to NP (such a result is stated explicitly by Leivant [Le]). Thus, if Kolaitis and Vardi's result could be extended to show that

[2]Technically, $\hat{S}$ is not a function, since its domain is not a set. If this bothers us, we can restrict our attention to universes that are, say, sets of natural numbers. Or, as is commonly done, we could assume that if the universe is of size n, then the universe is $\{0, 1, \ldots, n-1\}$.

strict Σ_1^1 *with successor* is not closed under complement, then this is equivalent to showing that NP is not closed under complement!

We now mention how we extend our results to allow more general built-in relations than successor, and comment as to why this is of interest. Intuitively, we allow built-in relations with no short cycles (when we ignore the directions of edges) and with small indegree and outdegree for each vertex. We assume that l (the number of built-in relations) is fixed, and we let $\hat{B}_i$ (for $i = 1, \dots, l$) be functions with domain the collection of possible universes of directed graphs, and range the collection of binary relations, such that the universe of the binary relation $\hat{B}_i(V)$ is V. We let $\Gamma(V)$ be the undirected graph with universe V where if v and w are distinct vertices, then (v, w) is an edge of $\Gamma(V)$ iff $\langle v, w\rangle \in \hat{B}_i(V)$ or $\langle w, v\rangle \in \hat{B}_i(V)$ for some i (where $1 \leq i \leq l$). Note that we are using the convention that $\langle v, w\rangle$ represents a directed edge, and (v, w) an undirected edge. We denote the cardinality of V by $|V|$.

THEOREM 3.2. *Let l (the number of built-in relations) be fixed. Assume that $\xi(n) \to \infty$ and $\sigma(n) \to 0$ as $n \to \infty$ (where $\sigma(n) > 0$ for every n). Assume also that $\Gamma(V)$ contains no cycle of length less than $\xi(|V|)$, and the degree of each point in $\Gamma(V)$ is at most $|V|^{\sigma(|V|)}$. Then directed reachability is not monadic Σ_1^1, even in the presence of the built-in relations.*

Thus, Theorem 3.2 says that there is no monadic Σ_1^1 sentence $\varphi(P, s, t, B_1, \dots, B_l)$ and functions $\hat{B}_1, \dots, \hat{B}_l$, with the restrictions given, such that $G = \langle V, P, s, t\rangle$ is (s, t)-connected iff $\langle V, P, s, t, \hat{B}_1(V), \dots, \hat{B}_l(V)\rangle \models \varphi$.

Similarly to before, we can consider an alternate notion of "in the presence of the built-in relations". Using this notion, "directed reachability is monadic Σ_1^1 in the presence of the built-in relations" would mean that there is a monadic Σ_1^1 sentence $\varphi(P, s, t, B_1, \dots, B_l)$ such that if $\langle V, P, s, t, B_1, \dots, B_l\rangle$ is a graph along with $B_1, \dots, B_l$ that are restricted as above, then $\langle V, P, s, t\rangle$ is (s, t)-connected iff $\langle V, P, s, t, B_1, \dots, B_l\rangle \models \varphi$. Unlike the situation with successor alone, this notion is not equivalent to our notion. However, under our notion, the nonexpressibility result is stronger (that is, implies the nonexpressibility result under the alternate notion). This alternate notion does not seem very interesting. But our notion is interesting, since $B_1, \dots, B_l$ can be specially chosen to be as useful as possible. We note that our notion of "built-in relations" is much like similar notions in [BoSi] and [Im3], where functions analogous to our assignment functions $\hat{B}_i$ are used that take as their argument the size of the universe, rather than the universe itself (these papers have the convention that if the universe is of size n, then the universe is $0, 1, \dots, n-1$, and so there is only one universe of each size). The notion of "built-in relations" in [Aj] is the same as ours; there, universes are restricted to be sets of natural numbers.

Now in the successor relation, there are no cycles at all, and the degree of each point is at most two. Hence, Theorem 3.1 is a special case of Theorem 3.2, where (a) $l = 1$, (b) $\hat{B}_1(V)$ is an arbitrary successor relation over V, (c) $\xi(n) = n$, and (d) $\sigma(n) = \log_n 2$.

There is another reason (besides our interest in successor relations) to allow built-in relations. Proving that a class is not, say, monadic Σ_1^1 shows that the class cannot be captured in a certain *uniform* way, where we think of a (fixed, finite length) monadic Σ_1^1 sentence as a uniform description. Proving that a class is not monadic Σ_1^1 even in the presence of certain built-in relations shows that the class cannot even be captured in certain *nonuniform* ways (since the built-in relations vary from universe to universe). So allowing built-in relations makes our nonexpressibility result that much more powerful.

We do not know whether our restriction on the built-in relations as given in the statement of Theorem 3.2 is essential: we consider it possible that directed reachability is not monadic Σ_1^1, even in the presence of arbitrary built-in relations of arbitrary arity (sometimes called "a polynomial amount of advice"). For that matter, we also consider it possible that connectivity (directed or undirected) is not monadic Σ_1^1, even in the presence of arbitrary built-in relations of arbitrary arity. In particular, it is an interesting open problem as to whether either directed or undirected connectivity or directed reachability is monadic Σ_1^1 in the presence of a built-in linear order (in the case of connectivity, this problem was originally posed by de Rougemont [Ro]). Note that Theorem 3.2 does not apply, since the undirected version of a linear order on at least three points has a cycle of length three. Another interesting special case would be when addition is built-in as in [Ly] (this is the ternary relation on the universe $\{0,\dots,n-1\}$ consisting of tuples (i,j,k) where $i+j=k$). We note that Lynch's results in [Ly] on monadic Σ_1^1 in the presence of a built-in addition relation do not apply, since Lynch encodes graphs as binary strings; hence, a unary relation for him over a graph with n vertices would have n^2 "bits", and so would correspond to a binary relation for us. Yet another interesting special case would be with Immerman's BIT relation built-in [Im4], consisting of tuples (i,j) where the ith bit in the binary expansion of j is a one.

§4. Games. In this section, we consider Ehrenfeucht-Fraïssé-type games [Eh], [Fr], describe how they have been used before to prove nonexpressibility results, and conclude by discussing the novel way we make use of them to prove our main result and its extensions.

We begin with an informal definition of an *r-round Ehrenfeucht-Fraïssé game* (where r is a positive integer). There are two *players*, called the *spoiler* and the *duplicator*, and two graphs, G_0 and G_1. In the first round, the spoiler selects a point in one of the two graphs, and the duplicator selects a point in the other 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 graphs, and the duplicator selects a point in the other 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 subgraph of G_0 induced by $\langle a_1,\dots,a_r\rangle$ is isomorphic to the subgraph of G_1 induced by $\langle b_1,\dots,b_r\rangle$ (under the function that maps a_i onto b_i for $1 \le i \le r$). That is, for the duplicator to win, there must be an edge between a_i and a_j in G_0 iff there is an edge between b_i and b_j in G_1, for each i, j. 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.

We shall sometimes refer to an r-round Ehrenfeucht-Fraïssé game as a *first-order game*. The following important theorem (from [Eh] and [Fr]) shows why these games are of interest.

THEOREM 4.1. *A class $\mathscr{S}$ of graphs is first-order definable iff there is r such that whenever $G_0 \in \mathscr{S}$ and $G_1 \notin \mathscr{S}$, then the spoiler has a winning strategy in the first-order game with parameters G_0, G_1, r.*

We now explain why the "only if" direction of Theorem 4.1 holds. Let φ be a first-order sentence which, when written in prenex normal form, is of the form $Q_1x_1 \cdots Q_rx_r\psi$, where Q_i is either $\forall$ or $\exists$ for each i, and where ψ is quantifier-free. Thus, when written in prenex normal form, φ has r quantifiers. Assume that $\mathscr{S}$ is the class of all graphs that satisfy φ. Assume also that $G_0 \in \mathscr{S}$ (so $G_0 \models \varphi$), and $G_1 \notin \mathscr{S}$ (so $G_1 \not\models \varphi$). We now show, by example, why the spoiler has a winning strategy in the first-order game with parameters G_0, G_1, r. Assume that, say, φ is $\forall x_1 \exists x_2 \psi(x, y)$, where ψ is quantifier-free. So $\exists x_1 \forall x_2 \neg \psi(x, y)$ is true about G_1. On the first round, the spoiler selects a point b_1 in G_1 such that $\forall x_2 \neg \psi(b_1, x_2)$ is true about G_1. Thus, intuitively, the spoiler makes "the existential move". Let a_1 be the duplicator's selection in G_0. Since $\forall x_1 \exists x_2 \psi(x, y)$ is true about G_0, it follows that $\exists x_2 \psi(a_1, y)$ is true about G_0. On the second round, the spoiler selects a point a_2 in G_0 such that $\psi(a_1, a_2)$ is true about G_0. Thus again, the spoiler "makes the existential move". Let b_2 be the duplicator's selection in G_1. So $\neg\psi(b_1, b_2)$ is true about G_1. Since $\psi(a_1, a_2)$ is true about G_0 and $\neg\psi(b_1, b_2)$ is true about G_1, and since ψ is quantifier-free, it follows easily that the subgraph of G_0 induced by $\langle a_1, a_2 \rangle$ is not isomorphic to the subgraph of G_1 induced by $\langle b_1, b_2 \rangle$. So indeed, the spoiler has won.

The converse to Theorem 4.1 tells us that if $\mathscr{S}$ is *not* first-order definable, then, in principle, this can be proven by a game-theoretic argument (by showing that for every r, there are $G_0 \in \mathscr{S}$ and $G_1 \notin \mathscr{S}$ such that the duplicator has a winning strategy in the first-order game with parameters G_0, G_1, r).

In addition to considering first-order games where the parameters are graphs G_0 and G_1 and a positive integer r, it is also convenient to consider first-order games where the parameters are a class $\mathscr{S}$, along with r. The rules of this game are as follows. The duplicator begins by selecting a member of $\mathscr{S}$ to be G_0, and a member of $\bar{\mathscr{S}}$ (the complement of $\mathscr{S}$) to be G_1. The players then play an r-round Ehrenfeucht-Fraïssé game. The duplicator wins precisely if he wins the r-round Ehrenfeucht-Fraïssé game. The next theorem follows easily from Theorem 4.1.

THEOREM 4.2. *$\mathscr{S}$ is first-order definable iff there is r such that the spoiler has a winning strategy in the first-order game with parameters $\mathscr{S}$, r.*

We now show how first-order games can be used to prove that evenness ("the number of points is even") is not expressible in first-order logic. Let $\mathscr{S}$ be the class of graphs with an even number of points, and let r be arbitrary. By Theorem 4.2, we need only show that the duplicator has a winning strategy in the first-order game with parameters $\mathscr{S}$, r. The duplicator begins by selecting G_0 to be a graph with $2r$ points and no edges, and G_1 to be a graph with $2r + 1$ points and no edges. In the r-round Ehrenfeucht-Fraïssé game that follows, the duplicator's strategy is simple. If on round i the spoiler selects a point that has not previously been selected by either player, then on round i the duplicator does the same in the other graph; if instead, the spoiler selects a point a_j (resp. b_j) that was previously selected by either player, then the duplicator selects b_j (resp. a_j).

We now discuss a modification of first-order games, called *monadic Σ_1^1 games*, introduced in [Fa2] to prove that connectivity is not monadic Σ_1^1 (although, as we saw, connectivity is monadic Π^1_1). The inputs are graphs G_0 and G_1, and positive integers c and r. 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-round Ehrenfeucht-Fraïssé game. The duplicator now wins if the subgraph of G_0 induced by $\langle a_1, \ldots, a_r \rangle$ is isomorphic to the subgraph of G_1 induced by $\langle b_1, \ldots, b_r \rangle$ (under the function that maps a_i onto b_i for $1 \leq i \leq r$), as before, and if in addition, the colors assigned to a_i and b_i are identical, for $1 \leq i \leq r$. That is, the isomorphism must also respect color.

As an example of such a game, assume that $c = r = 2$, that G_0 is the disjoint union of two cycles, and G_1 is a single cycle. We now show that the spoiler has a winning strategy (this corresponds to the fact that *non*connectivity is monadic Σ^1_1). Let us call the two colors red and blue. The spoiler's winning strategy is as follows. The spoiler colors all of the points in one cycle of G_0 red, and all of the points in the other cycle of G_0 blue. If the duplicator does not color some point of the single cycle G_1 red, then the spoiler can guarantee a win in round 1 by selecting a red point in G_0. So to have a chance to win, the duplicator must color some point of G_1 red, and similarly, the duplicator must color some point of G_1 blue. There is therefore some adjacent pair in G_1 with distinct colors; the spoiler selects these points in rounds 1 and 2 of the (first-order) Ehrenfeucht-Fraïssé game. The duplicator is forced to try to find an adjacent pair in G_0 with distinct colors, which is impossible.

The following theorem (from [Fa2]) is analogous to Theorem 4.1.

THEOREM 4.3. *A class $\mathscr{S}$ of graphs is monadic Σ^1_1 iff there are c and r such that whenever $G_0 \in \mathscr{S}$ and $G_1 \notin \mathscr{S}$, then the spoiler has a winning strategy in the first-order game with parameters G_0, G_1, c, and r.*

The "only if" direction of this theorem follows similarly to before. Thus, assume that φ is a monadic Σ^1_1 sentence $\exists A_1 \cdots \exists A_k \psi$, that G_0 is a graph that satisfies φ, and that G_1 is a graph that does not satisfy φ. Let $c = 2^k$, and let r be the number of quantifiers of the prenex normal form of ψ. As before, it is straightforward to show that "making the existential move" gives the spoiler a winning strategy in the monadic Σ^1_1 game with parameters G_0, G_1, c, r (the 2^k colors represent the various possibilities as to which of the sets corresponding to $A_1, \ldots, A_k$ a point is a member of).

Similarly to before, we wish to consider monadic Σ^1_1 games where the parameters are a class $\mathscr{S}$, along with positive integers c and r. The rules of this game are as follows.

1. The duplicator selects a member of $\mathscr{S}$ to be G_0.
2. The duplicator selects a member of $\overline{\mathscr{S}}$ 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-round Ehrenfeucht-Fraïssé game.

The winner is decided as before. The next theorem follows easily from Theorem 4.3.

THEOREM 4.4. *$\mathscr{S}$ is monadic Σ^1_1 iff there are c and r such that the spoiler has a winning strategy in the monadic Σ^1_1 game with parameters $\mathscr{S}$, c, r.*

In [Fa2] it is shown that, given c and r, there is a graph G_0 that is a large cycle, and a graph G_1 that is the disjoint union of two large cycles, such that the duplicator has a winning strategy in the monadic Σ_1^1 game with parameters G_0, G_1, c, r. Since G_0 is connected and G_1 is not, it follows (from Theorem 4.3 or Theorem 4.4) that connectivity is not monadic Σ_1^1. It is interesting to note that, as our example above showed, if we were to exchange G_0 for G_1 and vice-versa (so that G_0 were the disjoint union of two cycles, and G_1 the single cycle), then it would be the spoiler, rather than the duplicator, who has the winning strategy (provided c and r are at least 2).

The way that de Rougemont [Ro] extends [Fa2] to prove that connectivity is not monadic Σ_1^1, even in the presence of a built-in successor relation, is to define for each pair c, r a pair G_0, G_1 of graphs (that are substantially more complicated than disjoint unions of cycles, but still of bounded degree), define a successor relation for each, and then show that the duplicator has a winning strategy in the monadic Σ_1^1 game with parameters G_0, G_1, c, r, where the isomorphism must respect not only the graph relations and the coloring, but also the successor relations.

We do not see how to use Theorem 4.4 to prove our main result that directed reachability is not monadic Σ_1^1. Thus, when $\mathscr{S}$ is the class of directed graphs that are (s, t)-connected, and given c and r, we do not know opening moves for the duplicator (how to select G_0 and G_1) in a winning strategy for the monadic Σ_1^1 game. We therefore consider another game, the *new monadic Σ_1^1 game*, which, on the face of it, is easier for the duplicator to win. The inputs are a class $\mathscr{S}$, along with positive integers c and r. The rules of the new game are obtained from the rules of the monadic Σ_1^1 game by reversing the order of two of the moves. Specifically, the rules of the new monadic Σ_1^1 game are as follows.

1. The duplicator selects a member of $\mathscr{S}$ to be G_0.
2. The spoiler colors G_0 with the c colors.
3. The duplicator selects a member of $\bar{\mathscr{S}}$ to be G_1.
4. The duplicator colors G_1 with the c colors.
5. The spoiler and duplicator play an r-round Ehrenfeucht-Fraïssé game.

The winner is decided as before. Thus, in the new 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 new monadic Σ_1^1 game, we may sometimes refer to the monadic Σ_1^1 game as the *old* monadic Σ_1^1 game. In spite of the fact that it seems to be harder for the spoiler to win the new monadic Σ_1^1 game than the old monadic Σ_1^1 game, we have the following analogue to Theorem 4.4.

THEOREM 4.5. *$\mathscr{S}$ is monadic Σ_1^1 iff there are c and r such that the spoiler has a winning strategy in the new monadic Σ_1^1 game with parameters $\mathscr{S}$, c, r.*

The reason that Theorem 4.5 follows from Theorem 4.4 is that if $\mathscr{S}$ is monadic Σ_1^1, then the spoiler does not need to know what G_1 is before coloring G_0, since the coloring of G_0 is determined completely by "making the existential moves" as given by the monadic Σ_1^1 sentence that defines $\mathscr{S}$. The simple details are left to the reader.

There does not seem to be a new monadic Σ_1^1 game with parameters G_0, G_1, c, r that corresponds to the new monadic Σ_1^1 game with parameters $\mathscr{S}$, c, r in the same way as the old monadic Σ_1^1 game with parameters G_0, G_1, c, r corresponds to the old monadic Σ_1^1 game with parameters $\mathscr{S}$, c, r. Intuitively, this is because once we fix G_0 and G_1, it does not make sense to consider the notion that the spoiler must color G_0

without knowing what G_1 is. In fact, the reason that we bothered to introduce the first-order game and the old monadic Σ_1^1 game with $\mathscr{S}$ rather than G_0 and G_1 as (some of the) parameters is to contrast them with the new monadic Σ_1^1 game, which requires $\mathscr{S}$ as a parameter.

Unfortunately, even in the new monadic Σ_1^1 game, we do not know opening moves for the duplicator (how to select G_0 and G_1) in a winning strategy. However, we can prove that such moves exist! We do so by working in the spirit of the well-known probabilistic approach, where to prove that some object obeys property $\mathscr{P}$, we define a probability distribution on the set of objects, and show that nearly all objects (under this probability distribution) obey property $\mathscr{P}$. Of course, all that is needed is that the probability of property $\mathscr{P}$ is positive, but it is often convenient to prove that the probability of property $\mathscr{P}$ is nearly 1.

Let c and r be given. We construct (by probabilistic methods) a directed graph G with points s, t where there is a directed path from s to t in G. Thus, G is (s, t)-connected. As before, denote the graph that is obtained by deleting the edge e from G by $G - e$. We show that however the spoiler colors G with the c colors, there is an edge e of G (randomly selected from the set of edges of a certain type) such that (a) there is no directed path from s to t in $G - e$, and (b) when $G - e$ is colored in precisely the same way, vertex for vertex, as G was colored, then with positive probability (in fact, with high probability), the duplicator has a winning strategy in the r-round Ehrenfeucht-Fraïssé game played on G and $G - e$ (where, as before, the isomorphism must also respect color). Thus, in the new monadic Σ_1^1 game, the duplicator first selects G as G_0, then the spoiler colors G_0, then the duplicator selects G_e as G_1, then the duplicator colors G_1 by mimicking the coloring of G_0, and finally the duplicator wins the r-round Ehrenfeucht-Fraïssé game played on G_0 and G_1. It follows from Theorem 4.5 that directed connectivity is not monadic Σ_1^1.

There are several ways that this "game argument" differs from usual Ehrenfeucht-Fraïssé game arguments.

(1) The first difference corresponds to the difference between the old and new versions of the monadic Σ_1^1 game. In the old game, graphs G_0 and G_1 are specified before the spoiler colors G_0. Intuitively, the spoiler "knows" what G_0 and G_1 are before he colors G_0. In our case, this might be devastating for the duplicator: if the spoiler knew which edge e were deleted from $G_0 = G$ 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 our 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.

(2) Another difference is that in the old monadic Σ_1^1 game, the "hard part" of the duplicator's strategy is to find a coloring for G_1, once the duplicator knows how the spoiler has colored G_0. In this case, this is the easy part: the duplicator simply copies, vertex for vertex, the coloring of G_0.

(3) Our construction is probabilistic, rather than deterministic. In fact, as we noted, we do not know how to specify explicitly graphs G_0 and G_1 where our arguments work; we simply prove, by probabilistic arguments, that they exist. We come back to related issues at the end of this section.

We now describe in detail our construction of G and $G - e$. Assume that V is a finite set with n points. (There may also be built-in binary relations $B_1, \ldots, B_l$ with universe V.) Let p be a real number between 0 and 1. We now define a random directed graph Q_p^n on V. We refer to this graph as a *random path with random backedges*. Let $<$ be a random strict linear ordering of V with uniform distribution on the set of all strict linear orderings. If w is the immediate successor of v in the linear ordering $<$, then $\langle v, w \rangle$ is a *forward edge*. For each pair $v < w$ of distinct points, there is a *backedge* $\langle w, v \rangle$ with probability p. Each choice of backedges is made independently. Note that the presence of a backedge involving v and w (that is, a backedge $\langle v, w \rangle$ or $\langle w, v \rangle$, as appropriate) is independent of the choice of random linear order. The edges of Q_p^n are precisely the forward edges and the backedges. The least element of the ordering $<$ is s and the greatest is t. Clearly s and t are connected by a directed path in Q_p^n. Also, if we delete a forward edge e of the graph, then there is no longer a directed path from s to t, since the backedges all go from larger points in the linear ordering to smaller points. Thus, $Q_p^n - e$ is not (s, t)-connected.

Let L be a language containing $l + 1$ binary relation symbols (to represent the binary relation Q_p^n and the built-in binary relations $B_1, \ldots, B_l$), and constant symbols to represent s and t. Let $B_1, \ldots, B_l$ be binary relations with universe V. Define an undirected graph Γ as in §3: if v and w are distinct vertices, let (v, w) be an edge of Γ iff $\langle v, w \rangle \in B_i$ or $\langle w, v \rangle \in B_i$ for some i (where $1 \le i \le l$). We say that $\langle B_1, \ldots, B_l \rangle$ is a *successor-like family* (*with parameters* ξ *and* σ) if Γ contains no cycle of length less than ξ, and the degree of each point in Γ is at most $|V|^\sigma$.

To show that (s, t)-connectivity is not monadic Σ_1^1, it is enough to show that for each monadic Σ_1^1 sentence φ, if $\varepsilon > 0$ and if n is sufficiently large, then, for a suitable p, with probability at least $1 - \varepsilon$ the following holds: there is a forward edge e in Q_p^n so that if $Q_p^n \models \varphi$ then $(Q_p^n - e) \models \varphi$. In fact, we even allow a successor-like family of built-in relations. Let $\mathcal{M}$ be the *structure* (or *L-structure*) $\langle V, Q_p^n, s, t, B_1, \ldots, B_l \rangle$, and let $\mathcal{M}_e$ be the result of replacing Q_p^n in $\mathcal{M}$ by $Q_p^n - e$. We show that under certain assumptions, if φ is a monadic Σ_1^1 L-sentence (that is, φ is allowed to refer also to the built-in relations $B_1, \ldots, B_l$) and if $\langle B_1, \ldots, B_l \rangle$ is a successor-like family, then with high probability there is a forward edge e in Q_p^n so that if $\mathcal{M} \models \varphi$ then $\mathcal{M}_e \models \varphi$. As we noted, to prove our nonexpressibility results, we do not need "with high probability"; we could get by with "with positive probability". But we do prove the stronger result. Specifically, we prove the following theorem.

THEOREM 4.6. *Let r and l be positive integers, and let $\varepsilon > 0$. Assume that ξ is sufficiently large with respect to ε, that θ is sufficiently large with respect to r, l, and ε, that $\sigma > 0$ is sufficiently small with respect to ξ, and that n is sufficiently large with respect to r, l, ε, ξ, and σ. Suppose that V is a set with n elements, and that $\langle B_1, \ldots, B_l \rangle$ is a successor-like family with parameters ξ and σ. Assume that $\theta/n < p < n^\sigma/n$, that Q_p^n is a random path on V with random backedges, and that e is a random forward edge (with uniform distribution on the set of all forward edges). Then for each monadic Σ_1^1 sentence φ of the language L and of length at most r, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n and the deleted forward edge e together) that if $\mathcal{M} \models \varphi$ then $\mathcal{M}_e \models \varphi$.*

Note 1. Instead of saying "of length at most r", we could just as well have said "where the number of quantifiers in the prenex normal form of the first-order part of φ is at most r", but we phrased it as we did for ease in description.

Note 2. We cannot replace "if $\mathcal{M} \models \varphi$ then $\mathcal{M}_e \models \varphi$" in the conclusion of Theorem 4.6 by "$\mathcal{M} \models \varphi$ iff $\mathcal{M}_e \models \varphi$". This is because, as we noted, there is a monadic

Σ_1^1 sentence that says "The graph is not (s,t)-connected". Then $\mathcal{M}_e \models \varphi$ for every e, but $\mathcal{M} \not\models \varphi$.

Note 3. It is interesting that the probability p in the random construction for Q_p^n cannot be too small or too large. If $e = \langle x, y \rangle$ is the deleted forward edge, we need p to be large enough that there are enough backedges from x to "play the role" of the missing edge e. On the other hand, we need p to be small enough that the expected degree of each vertex in Q_p^n is small (just as we want the degree of each vertex in Γ to be small), so that with high probability, the number of points in short cycles (when we ignore the directions of edges) is small (this turns out to be important in our proof that with high probability, the duplicator has a winning strategy). In fact, we originally tried to prove our result by taking $p = 1/2$; instead we proved, to our surprise, that if p is constant, then even if there are no built-in relations, with high probability it is the spoiler, rather than the duplicator, who has the winning strategy! We prove this in §6. Thus, if p is too small (such as $p = 0$ or p very near 0 as a function of the size n of the universe) or too large (such as p a positive constant) then with high probability the spoiler has a winning strategy; if p is in an intermediate range, then with high probability the duplicator has a winning strategy.

The proof of Theorem 4.6 appears in §8. If we ignore the built-in relations $B_1, \ldots, B_l$, then Theorem 4.6 implies our main result that directed reachability is not monadic Σ_1^1. For, assume that φ were a Σ_1^1 sentence that characterizes directed reachability. By Theorem 4.6, there are $\mathcal{M}$ and $\mathcal{M}_e$ such that if $\mathcal{M} \models \varphi$ then $\mathcal{M}_e \models \varphi$. But this is impossible, since $\mathcal{M}$ is (s,t)-connected although $\mathcal{M}_e$ is not. If we consider the built-in relations, then Theorem 4.6 implies Theorem 3.2 above, as we now show. Assume that φ were a Σ_1^1 sentence of length r that characterizes directed reachability in the presence of the l built-in relations. By Theorem 4.6, we can find ξ_0, σ_0, and n_0 so that if $n \geq n_0$, then for each successor-like family $\langle B_1, \ldots, B_l \rangle$ with parameters ξ_0 and σ_0, there is $\mathcal{M}_e$ (with universe of size n) that satisfies φ. Find $n_1 \geq n_0$ so that $\xi(n) > \xi_0$ and $0 < \sigma(n) < \sigma_0$ whenever $n \geq n_1$. Let V be a universe of size n_1. Then $\langle \hat{B}_1(V), \ldots, \hat{B}_l(V) \rangle$ is a successor-like family with parameters $\xi(n_1)$ and $\sigma(n_1)$. Since $\xi(n_1) > \xi_0$ and $\sigma(n_1) < \sigma_0$, it follows that $\langle \hat{B}_1(V), \ldots, \hat{B}_l(V) \rangle$ is a successor-like family with parameters ξ_0 and σ_0. Hence, $\mathcal{M}_e$ satisfies φ, a contradiction, since $\mathcal{M}_e$ is not (s,t)-connected.

We close this section by giving an interesting twist on the point made above that we do not know opening moves for the duplicator in a winning strategy for the old monadic Σ_1^1 game. That is, given c and r, we do not know how to construct a pair G_0', G_1' of directed graphs where G_0' is (s,t)-connected, G_1' is not (s,t)-connected, and the duplicator has a winning strategy in the monadic Σ_1^1 game with parameters G_0', G_1', c, r. Since directed reachability is not monadic Σ_1^1, it follows from Theorem 4.4 that for each pair c, r, there is such a pair G_0', G_1'. In our proof of non-expressibility, we work 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 removing some forward edge from G_0. The interesting point is that it is not clear that such a pair G_0, G_1 could serve as G_0', G_1'! From Theorem 1.2, we do know that as c and r grow, the indegree or outdegree of G_0' or G_1' must grow arbitrarily large. In particular, graphs of the type used in [Fa2] and [Ro] are inadequate, since they have bounded degree.

§5. The bounded-degree case. In this section, we give a proof of the following result:

THEOREM 5.1. *Let k be a positive integer. The class of directed (s,t)-connected graphs where the indegree and outdegree of each vertex is at most k is monadic Σ_1^1.*

PROOF. Let φ be the Σ_1^1 sentence $\exists A_1 B_1 A_2 B_2 \cdots A_{2k^2-1} B_{2k^2-1} \psi$, where the A_i's and B_i's are unary relation symbols, and where ψ is a first-order formula that we now describe informally.

Assume that G is the graph in question, and $V(G)$ (resp. $E(G)$) is its set of vertices (edges). If A_i and B_i are subsets of $V(G)$, then we may define a subset D_i of $E(G)$ by saying $\langle x, y\rangle \in D_i$ iff $x \in A_i \wedge y \in B_i \wedge \langle x, y\rangle \in E(G)$. Let $D = \bigcup_{i=1}^{2k^2-1} D_i$. Then ψ is a first-order formula that "says" that the indegree and outdegree of each vertex is at most k, along with the following conditions:

1. s is the tail of exactly one edge in D, and s is not the head of any edge in D.[3]
2. t is the head of exactly one edge in D, and t is not the tail of any edge in D.
3. If $x \in A_i$ or $x \in B_i$, and if $x \neq s$ and $x \neq t$, then x is the head of exactly one edge in D and x is the tail of exactly one edge in D.

First we prove that φ implies (s,t)-connectivity. Indeed, if φ holds then we start a path from s consisting of edges in D. Our conditions imply that this path cannot go through a vertex twice, and cannot end anywhere else than t.

To show that (s,t)-connectivity (along with the indegree and outdegree of each vertex being at most k) implies φ, it is sufficient to show that there are subsets A_i and B_i, for $1 \le i \le 2k^2 - 1$, so that if we define the set D in the way just described, it satisfies the three conditions.

Let $s = s_1, \ldots, s_l = t$ be a minimal path from s to t, and let f_j be the edge $\langle s_j, s_{j+1}\rangle$, for $1 \le j < l$. We will define the sets A_i and B_i (for $1 \le i \le 2k^2 - 1$) so that D will be precisely the set of edges f_j, for $1 \le j \le l - 1$. Clearly this implies the three conditions.

First we define an undirected graph H whose vertices consist of the set $E(G)$ of edges of G. If $e_1 = \langle x_1, y_1\rangle \in E(G)$ and $e_2 = \langle x_2, y_2\rangle \in E(G)$, then (e_1, e_2) is an edge of H precisely if $e_1 \neq e_2$ and either $\langle x_1, y_2\rangle \in E(G)$ or $\langle x_2, y_1\rangle \in E(G)$. Thus, (e_1, e_2) is an edge of H, where $e_1 \neq e_2$ precisely if there is an edge e with the same tail as one of e_1 or e_2 and with the same head as the other of e_1 or e_2. Because indegrees and outdegrees in G are at most k, it is straightforward to verify that the degree of each vertex of H is at most $2(k^2 - 1)$. This implies that H can be colored with $2(k^2 - 1) + 1 = 2k^2 - 1$ colors so that neighboring points in H are of different colors (we simply use a greedy coloring algorithm). For $1 \le i \le 2k^2 - 1$, let C_i be the set of points in H with the ith color, and let A_i (resp. B_i) be the set of all $x \in V(G)$ such that x is the tail (the head) of some edge f_j as defined above, where $f_j \in C_i$.

As we noted, the proof is complete if we show that D consists precisely of the edges f_j (for $1 \le j \le l - 1$). Clearly, each f_j is in D. Assume now that $e = \langle x, y\rangle \in D$. Then there is an i so that the tail of e is in A_i and the head of e is in B_i. The definitions of A_i and B_i imply that $x = s_j$ and $y = s_m$ for some j and m. Since $s_j \in A_i$, this implies that $f_j \in C_i$. Similarly, since $s_m \in B_i$, this implies that $f_m \in C_i$. Since e is an edge with the same tail as f_j and the same head as f_m, it follows that either $f_j = f_m$ or else (f_j, f_m) is an edge of H. But (f_j, f_m) cannot be an edge of H, since f_j and f_m have the same color (namely, the ith color). So $f_j = f_m$. Since e has the same tail as f_j and the same head as $f_m = f_j$, it follows that $e = f_j$, as desired. ■

[3]The *tail* of an edge $\langle v, w\rangle$ is v, and the *head* is w.

§6. Constant probability of a backedge. In this section, we show that if the probability p of a backedge (as defined in §4) is constant (independent of n), then it is the spoiler, rather than the duplicator, who has the winning strategy with high probability. We begin by considering the simplest cases, namely, $p = 0$ and $p = 1$.

If $p = 0$ (that is, if there are no backedges), then the spoiler has a winning strategy with probability 1 in the 2-round Ehrenfeucht-Fraïssé game, even with no colors at all. This is because in Q_p^n there is only one point (namely, the minimal point s) with indegree 0, but in $Q_p^n - e$ there are two such points, namely s and the head of the deleted forward edge e.

If $p = 1$ (that is, if every possible backedge appears), then we now show that the spoiler has a winning strategy with probability 1 in the 4-round Ehrenfeucht-Fraïssé game, again with no colors at all. Throughout this section, whenever we speak of the distance $d(u, v)$ between two vertices u and v, we mean the distance with respect to the linear ordering $<$ (that is, the distance in Q_p^n where the backedges are completely ignored). Let $e = \langle b_1, b_2 \rangle$ be the deleted forward edge. On the spoiler's first two moves, he picks b_1 and b_2 in $Q_p^n - e$. Let a_1 and a_2 be the corresponding points selected by the duplicator in Q_p^n. Since $\langle b_2, b_1 \rangle$ is an edge of $Q_p^n - e$ but $\langle b_1, b_2 \rangle$ is not, we know that $\langle a_2, a_1 \rangle$ must be an edge of Q_p^n and $\langle a_1, a_2 \rangle$ must not be, or else the spoiler can guarantee a win after round 2. Hence $a_1 < a_2$, and $d(a_1, a_2) \geq 2$ (that is, a_2 is not the immediate successor of a_1). There are three cases.

(a) Assume first that $d(a_1, a_2) = 2$. Then on round 3 the spoiler selects the point a_3 in Q_p^n that is the immediate successor of a_1 (and the immediate predecessor of a_2). Then $\langle a_1, a_3 \rangle$, $\langle a_3, a_1 \rangle$, $\langle a_3, a_2 \rangle$, $\langle a_2, a_3 \rangle$ are all edges of Q_p^n. But there is no point b_3 in $Q_p^n - e$ such that $\langle b_1, b_3 \rangle$, $\langle b_3, b_1 \rangle$, $\langle b_3, b_2 \rangle$, $\langle b_2, b_3 \rangle$ are all edges of $Q_p^n - e$. So the spoiler has guaranteed a win after round 3.

(b) Assume now that $d(a_1, a_2) = 3$. Then on round 3 the spoiler selects the points a_3, a_4 in Q_p^n such that a_3 is the immediate successor of a_1 and such that a_4 is the immediate successor of a_3 (and the immediate predecessor of a_2). Then $\langle a_1, a_3 \rangle$, $\langle a_3, a_1 \rangle$, $\langle a_3, a_4 \rangle$, $\langle a_4, a_3 \rangle$, $\langle a_4, a_2 \rangle$, $\langle a_2, a_4 \rangle$ are all edges of Q_p^n. But there are no points b_3, b_4 in $Q_p^n - e$ such that $\langle b_1, b_3 \rangle$, $\langle b_3, b_1 \rangle$, $\langle b_3, b_4 \rangle$, $\langle b_4, b_3 \rangle$, $\langle b_4, b_2 \rangle$, $\langle b_2, b_4 \rangle$ are all edges of $Q_p^n - e$. So the spoiler has guaranteed a win after round 4.

(c) Assume finally that $d(a_1, a_2) \geq 4$. Then on round 3 the spoiler selects the point a_3 in Q_p^n such that a_3 is the immediate successor of the immediate successor of a_1. So $a_1 < a_3 < a_2$, and also $d(a_1, a_3) \geq 2$ and $d(a_3, a_2) \geq 2$. So $\langle a_2, a_3 \rangle$ and $\langle a_3, a_1 \rangle$ are edges of Q_p^n, but $\langle a_3, a_2 \rangle$ and $\langle a_1, a_3 \rangle$ are not. But there is no point b_3 in $Q_p^n - e$ such that $\langle b_2, b_3 \rangle$ and $\langle b_3, b_1 \rangle$ are edges of $Q_p^n - e$, and $\langle b_3, b_2 \rangle$ and $\langle b_1, b_3 \rangle$ are not. So the spoiler has guaranteed a win after round 3.

We now consider the more complicated case where $0 < p < 1$.

THEOREM 6.1. *Assume that $\varepsilon > 0$, that p is constant $(0 < p < 1)$, and that n is sufficiently large with respect to ε and p. Then, with probability at least $1 - \varepsilon$, there is a vertex coloring of Q_p^n (where the number of colors depends on p) such that if G is the resulting colored graph and e is a random forward edge (with uniform distribution on the set of all forward edges), then the spoiler has a winning strategy in the 6-round Ehrenfeucht-Fraïssé game played on G and $G - e$ (where isomorphisms must respect colors).*

PROOF. Let $M \geq 2$ and $N \geq 2$ be positive integers that we shall select later to be sufficiently large with respect to p. We now define a coloring with $3MN$ colors

$\Omega_1, \ldots, \Omega_{3MN}$. Let $m = \lfloor(\log n)/M\rfloor$. (For definiteness, let us take all logarithms to the base 2.) Let L_i be the ith interval of Q_p^n of size m, for $1 \le i \le \lfloor n/m \rfloor$. Thus, L_1 contains the first m points in the linear order, L_2 contains the next m points in the linear order, and so on. If m does not divide n, then define a final interval $L_{\lceil n/m \rceil}$ to contain the remaining points (this interval will contain less than m points). Each point in L_i has color Ω_j, where $1 \le j \le 3MN$ and $j \equiv i \bmod 3MN$. In particular, if u and v have the same color, then either u and v are in the same interval L_i for some i, or else

$$d(u,v) > (3MN - 1)m = (3MN - 1)\lfloor(\log n)/M\rfloor,$$

which is greater than $2N \log n$ if n is sufficiently large with respect to N.

Call L_i and L_{i+1} *adjacent intervals*. Also, call the colors Ω_i and Ω_{i+1} *adjacent colors* (where we take the subscript mod $3MN$). Thus, Ω_1 and Ω_2 are adjacent colors, as are Ω_1 and Ω_{3MN}. We define a ternary relation T over the universe V of Q_p^n, by letting $Tuwv$ hold iff the color of w is not the same as nor adjacent to either the color of u or the color of v, and if $\langle v, w\rangle$ and $\langle w, u\rangle$ are edges of Q_p^n. Because of the constraint on the color of w, it is clear that $Tuwv$ implies that $\langle v, w\rangle$ and $\langle w, u\rangle$ are both backedges. In particular, $Tuwv$ implies that $u < w < v$, and u and v are not in the same or even adjacent intervals. Let R be the ternary relation where $Ruwv$ holds iff either $Tuwv$ or $Tvwu$. In particular, $Ruwv$ implies that u and v are not in the same or even adjacent intervals.

We now show that if N is sufficiently large with respect to p, and if n is sufficiently large with respect to ε, then, with probability at least $1 - \varepsilon/4$, for every pair u, v of vertices where $d(u,v) \ge N \log n$, there is w such that $Ruwv$. Assume for now that u and v are fixed, that $u < v$, and that $d(u,v) \ge N \log n$. Call w a *u, v-candidate* if (a) $u < w < v$ and (b) the color of w is not the same as nor adjacent to either the color of u or the color of v. Call w a *u, v-winner* if w is a u, v-candidate and if $\langle v, w\rangle$ and $\langle w, u\rangle$ are both backedges. In particular, if w is a u, v-winner, then $Ruwv$ holds. It is easy to see that if N is sufficiently large, then there are at least $(N \log n)/2$ distinct u, v-candidates. For each u, v-candidate w, the probability that w is a u, v-winner is p^2, independently of the probability that other u, v-candidates are winners. Hence, the probability that there is no u, v-winner is at most

$$(1 - p^2)^{(N \log n)/2} = n^{(N/2)\log(1 - p^2)}.$$

So the probability that there is some u, v where $u < v$ and $d(u,v) \ge N \log n$ and there is no u, v-winner is less than $n^2 n^{(N/2)\log(1-p^2)}$. Assume that N is sufficiently large that $(N/2)\log(1 - p^2) < -3$ (of course, $\log(1 - p^2)$ is negative). Then this probability is less than $1/n$, which is less than $\varepsilon/4$ if n is sufficiently large with respect to ε. A similar result holds if $v < u$ and $d(u,v) \ge N \log n$. So assume that whenever $d(u,v) \ge N \log n$, there is w such that $Ruwv$.

Assume now that $e = \langle x, y\rangle$ is a randomly deleted forward edge. With probability $1/m = 1/\lfloor(\log n)/M\rfloor$, it happens that x and y are in distinct intervals L_i (that is, the edge e "straddles the border"). By taking n sufficiently large with respect to ε and M, this probability is less than $\varepsilon/4$. So assume that x and y are in the same interval L_i. Let L_i be $[b_1, b_2]$ (that is, the left endpoint is b_1, and the right endpoint is b_2). With probability at most

$$(\lceil n^{1/2}\rceil + m)/(n - 1) = (\lceil n^{1/2}\rceil + \lfloor(\log n)/M\rfloor)/(n - 1),$$

it happens that b_2 is among the final $\lceil n^{1/2} \rceil$ points in the linear order. By taking n sufficiently large with respect to ε, this probability is less than $\varepsilon/4$. So assume that b_2 is not among the final $\lceil n^{1/2} \rceil$ points in the linear order.

Let $q = \min\{p, 1 - p\}$. The probability that there is no point $b_3 > b_2$ such that there is a backedge from b_3 to every member of $[b_1, x]$, but no backedge from b_3 to any member of $[y, b_2]$, is at most

$$(1 - q^{(\log n)/M})^{n^{1/2}} = (1 - n^{(\log q)/M})^{n^{1/2}}.$$

We now select M to be sufficiently large that $(\log q)/M > -1/4$. This probability is then less than $(1 - n^{-1/4})^{n^{1/2}}$. If n is sufficiently large with respect to ε, then this probability is less than $\varepsilon/4$ (we can prove this by, say, l'Hospital's rule). So assume that there is a point $b_3 > b_2$ such that there is a backedge from b_3 to every member of $[b_1, x]$, but no backedge from b_3 to any member of $[y, b_2]$.

We now describe the spoiler's winning strategy. On rounds 1 and 2, the spoiler selects b_1, b_2, b_3 in $G - e$. Let a_1, a_2, a_3 be the corresponding points selected by the duplicator in G. Let us refer to the color of b_1 and b_2 as red. Then a_1 and a_2 are also red (or else the spoiler has guaranteed a win). We now show that if a_1 and a_2 are not in the same interval L_j for some j, then the spoiler can guarantee a win after round 4. For, assume that a_1 and a_2 are not in the same interval L_j for some j. Since a_1 and a_2 have the same color, we have shown that $d(a_1, a_2) > 2N \log n$. Hence, there is w such that Ra_1wa_2. The spoiler could then select $a_4 = w$ from G in round 4, and the duplicator would need to select b_4 from $G - e$ so that $Rb_1b_4b_2$. But this is impossible, since, as we showed, $Ruwv$ implies that u and v are not in the same interval. So we can assume that a_1 and a_2 are in the same interval L_j for some j.

We now show that if it is not the case that a_1 and a_2 are the left and right endpoints of this interval L_j (that is, if $L_j \neq [a_1, a_2]$), then the spoiler can guarantee a win after round 5. Assume that a_2 is not the right endpoint of the interval L_j (a similar argument works if a_1 is not the left endpoint of L_j). Let z be the immediate successor (under the linear ordering) of b_2. The spoiler could select z in $G - e$ as his move b_4 on round 4. Let the color of z be blue (it is different from red, the color of b_2, since b_2 is the right endpoint of its interval). In order to avoid losing, the duplicator must then select a blue point a_4 in G such that there is an edge $\langle a_2, a_4 \rangle$. Since a_2 is not the right endpoint of its interval, this edge $\langle a_2, a_4 \rangle$ would be a backedge. Since red and blue are adjacent colors, it follows that $d(a_2, a_4) \geq N \log n$. So there would be a_5 such that $Ra_2a_5a_4$. The spoiler could select a_5 in G on round 5, and the duplicator would then be forced to select b_5 in $G - e$ such that $Rb_2b_5b_4$. This is impossible, since, as we noted, $Ruwv$ implies that u and v are not in adjacent intervals. So $L_j = [a_1, a_2]$.

Let us refer to each point u where $\langle b_3, u \rangle$ is an edge of $G - e$ as a $+_b$ point; otherwise, we refer to u as a $-_b$ point. Similarly, let us refer to each point u where $\langle a_3, u \rangle$ is an edge of G as a $+_a$ point; otherwise, we refer to u as a $-_a$ point. Since b_1 is a $+_b$ point and b_2 is a $-_b$ point, it follows that a_1 must be a $+_a$ point and a_2 a $-_a$ point, or else the spoiler has guaranteed a win after round 3. Since there is a forward edge in G from each point to its successor in the linear ordering, it follows that there is some pair a_4, a_5 in $[a_1, a_2]$ such that a_4 is a $+_a$ point, a_5 is a $-_a$ point, and there is an edge $\langle a_4, a_5 \rangle$ in G. On rounds 4 and 5, the spoiler selects a_4 and a_5 in G. Let b_4 and b_5 be the duplicator's corresponding moves in $G - e$. Then b_4 is a

$+_b$ point and b_5 is a $-_b$ point, or else the spoiler has guaranteed a win after round 5.

We now show that if $b_4 \notin [b_1, b_2]$, then the spoiler has guaranteed a win after round 6. For, a_4 is red, since $a_4 \in [a_1, a_2]$. Hence, b_4 is red, or else the spoiler has guaranteed a win. If $b_4 \notin [b_1, b_2]$, then $d(b_1, b_4) > 2N \log n$, so there is w such that Rb_1wb_4. The spoiler could select w in $G - e$ as b_6 on round 6, and win as before, since there is no a_6 such that $Ra_1a_6a_4$, because a_1 and a_4 are in the same interval. Hence, $b_4 \in [b_1, b_2]$. Similarly, $b_5 \in [b_1, b_2]$. If the spoiler has not guaranteed a win after round 5, then $\langle b_4, b_5 \rangle$ is an edge of $G - e$, since $\langle a_4, a_5 \rangle$ is an edge of G. But then there is an edge in $G - e$ from a $+_b$ point in $[b_1, b_2]$ to a $-_b$ point in $[b_1, b_2]$. However, this is impossible, by construction. ∎

§7. **Large deviation theorem.** In this section, we give a result that we shall need to make use of several times. We begin with a theorem that is one variation of the well-known result that the probability that the number of successful independent trials differs from the mean by more than a constant times the mean is exponentially small. In this theorem "Pr" refers to the probability. The version of this theorem that we now state appears in [Bo, p. 13] (see also [Che]).

THEOREM 7.1 (Large Deviation Theorem). *Assume that there are n independent trials, each with probability p of success. Assume that* $0 < p \leq 1/2$, *that* $0 < c \leq 1/12$, *and that* $cp(1 - p)n \geq 12$. *Let* S_n *be the number of successful trials, and let* $M = pn$ *be the expected number of successful trials. Then*

$$\Pr\{|S_n - M| \geq cM\} \leq (c^2M)^{-1/2}e^{-c^2M/3}.$$

We shall make use several times of the following corollary to Theorem 7.1.

COROLLARY 7.2. *There are positive constants* c_1 *and* c_2 *such that the following holds. Let* $A_1, \ldots, A_n$ *be* 0, 1*-valued random variables, and let* S_n *be the cardinality of* $\{i \mid A_i = 1\}$. *Assume that* $0 < p \leq 1/2$, *and that* $M = pn \geq c_1$.

1. *Assume that for all i and for all* 0, 1*-sequences* $j_1, \ldots, j_{i-1}$, *the conditional probability* $\Pr\{A_i = 1 \mid A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}$ *is at most p (that is, the conditional probability that* $A_i = 1$, *given an arbitrary condition about* $A_1, \ldots, A_{i-1}$, *is at most p). Then* $\Pr\{S_n \geq 13M/12\} < e^{-c_2M}$.

2. *Assume that for all i and for all* 0, 1*-sequences* $j_1, \ldots, j_{i-1}$, *the conditional probability* $\Pr\{A_i = 1 \mid A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}$ *is at least p. Then* $\Pr\{S_n \leq 11M/12\} < e^{-c_2M}$.

PROOF. We just prove (1), since the proof of (2) is similar. We shall define independent random variables $B_1, \ldots, B_n$ such that, for each i (with $1 \leq i \leq n$), (a) $\Pr\{B_i = 1\} = p$, and (b) if $A_i = 1$ then $B_i = 1$. The result then follows from Theorem 7.1, where we let $c = 1/12$, $c_1 = 288$, and $c_2 = 1/432$.

For technical convenience, we assume that for each 0, 1-sequence $j_1, \ldots, j_n$, we have $\Pr\{A_1 = j_1, \ldots, A_n = j_n\} > 0$. If this is not the case, then we slightly modify the probabilities so that this is the case (by replacing each probability of 0 by a probability of ε, and appropriately lowering other probabilities slightly), prove the result, and then pass to the limit as $\varepsilon \to 0$ to prove the result in the original case where some joint probabilities may be 0.

Define $\eta_1, \ldots, \eta_n$ to be new random variables, all independent and independent of $A_1, \ldots, A_n$, and each uniformly distributed over [0, 1]. For each i with $1 \leq i \leq n$ and each 0, 1-sequence $j_1, \ldots, j_{i-1}$, define $p_{\langle j_1, \ldots, j_{i-1} \rangle}$ to be $\Pr\{A_i = 1 \mid A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}$. By assumption, $p_{\langle j_1, \ldots, j_{i-1} \rangle} \le p \le 1/2$.

For $1 \le i \le n$, define the new 0, 1-valued random variable B_i via

$$B_i = 1 \quad \text{iff} \quad A_i = 1 \text{ or}$$
$$\bigvee_{j_1, \ldots, j_{i-1}} \left(A_i = 0 \wedge A_1 = j_1 \wedge \cdots \wedge A_{i-1} = j_{i-1} \wedge \eta_i \le \frac{p - p_{\langle j_1, \ldots, j_{i-1} \rangle}}{1 - p_{\langle j_1, \ldots, j_{i-1} \rangle}} \right).$$

Clearly, if $A_i = 1$ then $B_i = 1$. Before we show that the B_i's are independent, and that $\Pr\{B_i = 1\} = p$, we first consider $\Pr\{B_i = 1 \mid A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}$. Denote $\langle j_1, \ldots, j_{i-1} \rangle$ by $\vec{j}$. Define $q_{\vec{j}}$ to be $\Pr\{A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}$. By assumption, $q_{\vec{j}} > 0$. We have

$$\begin{aligned}\Pr\{B_i = 1 \mid A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\} &= \frac{\Pr\{B_i = 1, A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}}{\Pr\{A_1 = j_1, \ldots, A_{i-1} = j_{i-1}\}} \\ &= \frac{q_{\vec{j}} p_{\vec{j}} + q_{\vec{j}}(1 - p_{\vec{j}})^{(p - p_{\vec{j}})/(1 - p_{\vec{j}})}}{q_{\vec{j}}} \\ &= p.\end{aligned}$$

Thus, $B_i = 1$ with probability p, independent of $A_1, \ldots, A_{i-1}$. Since $B_1, \ldots, B_{i-1}$ are defined in terms of $A_1, \ldots, A_{i-1}, \eta_1, \ldots, \eta_{i-1}$, and since the η_j's are independent of everything else, it follows that $B_i = 1$ with probability p, independent of $B_1, \ldots, B_{i-1}$. But then the B_i's are all independent, since

$$\begin{aligned}&\Pr\{B_1 = j_1, \ldots, B_n = j_n\} \\ &\quad = \Pr\{B_1 = j_1\}\Pr\{B_2 = j_2 \mid B_1 = j_1\} \cdots \Pr\{B_n = j_n \mid B_1 = j_1, \ldots, B_{n-1} = j_{n-1}\} \\ &\quad = \Pr\{B_1 = j_1\}\Pr\{B_2 = j_2\} \cdots \Pr\{B_n = j_n\}. \quad \blacksquare\end{aligned}$$

For convenience, in what follows, we shall always make use of Corollary 7.2 instead of Theorem 7.1, even in the case of independent trials, so that we can refer to the constants c_1 and c_2.

§8. Proof of Theorem 4.6. In this section, we prove Theorem 4.6, which as we have seen is sufficient to prove our nonexpressibility results. To prove this theorem, it is easy to see that it is sufficient to prove the following proposition.

PROPOSITION 8.1. *Let r and l be positive integers, and let $\varepsilon > 0$. Assume that ξ is sufficiently large with respect to ε, that θ is sufficiently large with respect to r, l, and ε, that $\sigma > 0$ is sufficiently small with respect to ξ, and that n is sufficiently large with respect to r, l, ε, ξ, and σ. Suppose that V is a set with n elements, and that $\langle B_1, \ldots, B_l \rangle$ is a successor-like family with parameters ξ and σ. Assume that $\theta/n < p < n^{\sigma}/n$, and that Q_p^n is a random path on V with random backedges. Let L' be the language obtained from L by including also symbols to represent unary relations $A_1, \ldots, A_k$, where $k \le r$. Let ψ be a first-order sentence of length at most r over the language L'.*

With probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), for each choice of the unary relations $A_1, \ldots, A_k$ over V, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that

$$\langle V, Q_p^n, s, t, B_1, \ldots, B_l, A_1, \ldots, A_k \rangle \models \psi$$
$$\textit{iff } \langle V, Q_p^n - e, s, t, B_1, \ldots, B_l, A_1, \ldots, A_k \rangle \models \psi.$$

Note. The reader may be puzzled as to why the symmetric conclusion ("iff") of Proposition 8.1 does not imply that we can replace "if $\mathscr{M} \models \varphi$ then $\mathscr{M}_e \models \varphi$" in the conclusion of Theorem 4.6 by "$\mathscr{M} \models \varphi$ iff $\mathscr{M}_e \models \varphi$", which we showed in Note 2 after Theorem 4.6 is not possible. The reason is that the sentence φ in that note is of the form $\exists A\psi$, where the choice of A depends on e (that is, no single set A "works" for every choice of e). Further details are left to the reader.

It is convenient to view $\langle V, Q_p^n, s, t, B_1, \ldots, B_l, A_1, \ldots, A_k\rangle$ as a *colored graph* G_0, which is a directed graph where there is a color associated with each vertex and each edge. The set $V(G_0)$ of vertices of G_0 is simply V. We assume that there are $2^k + 2$ vertex colors, where 2^k of the vertex colors tell precisely which of the sets $A_1, \ldots, A_k$ the vertex is a member of, and the remaining two colors tell whether or not the vertex is s and whether or not the vertex is t. The (directed) edges of G_0 consist of the union of the edges in Q_p^n, $B_1, \ldots, B_l$. We assume that there are 2^{l+1} edge colors, where the color of an edge tells precisely which of Q_p^n, $B_1, \ldots, B_l$ the edge is an edge of. When necessary, we may denote by C_0 the vertex and edge colorings we have given for G_0. Similarly, we let G_1 be the colored graph corresponding to $\langle V, Q_p^n - e, s, t, B_1, \ldots, B_l, A_1, \ldots, A_k\rangle$, and denote the coloring by C_1. We assume that the meaning of a color Ω in C_0 is the same as that of Ω in C_1. In particular, if $v \in V$, then the C_1 color of v is the same as the C_0 color of v; similarly, each edge other than e has the same C_1 color as its C_0 color.

To prove Proposition 8.1 (and hence our main theorem and its extension to allowing built-in relations), we need only prove the following result.

PROPOSITION 8.2. *Assume the first paragraph of assumptions in Proposition* 8.1. *With probability at least* $1 - \varepsilon$ *(where the probability is taken over the choices of* Q_p^n*), for each choice of the unary relations* $A_1, \ldots, A_k$ *over* V*, the probability is at least* $1 - \varepsilon$ *(where the probability is taken over the choices of the deleted forward edge* e*) that the duplicator has a winning strategy in the* r*-round Ehrenfeucht-Fraïssé game played on* G_0 *and* G_1*, where the isomorphisms must respect vertex and edge colors.*

To prove Proposition 8.2, we introduce vertex colorings that refine the vertex colorings of C_0 and C_1. If d and q are positive integers, we shall define the d, q-*coloring* of G_0 (the definition of the d, q-coloring of G_1 is the same, except that we replace G_0 and C_0 by G_1 and C_1 throughout). For each d, the $d, 1$-color of each vertex is simply the C_0 color. Assume inductively that the d, q-coloring has been defined; we now define the $d, q + 1$-coloring. Let v be a vertex. For each vertex w, define the v-*type* of w to be a complete description of the d, q-color of w, along with a complete description as to whether or not $\langle v, w\rangle$ is an edge of G_0, if so what its C_0 color is, whether or not $\langle w, v\rangle$ is an edge of G_0, and if so what its C_0 color is. The $d, q + 1$-color of v is a complete description of the d, q-color of v, along with a complete description, for each possible v-type, as to whether there are $0, 1, \ldots, d - 1$, or at least d points with that v-type. Thus, the $d, q + 1$-color of a vertex v tells the d, q-color of v, and also tells how many vertices there are of each v-type, except that we do not count beyond d.

It is important to note for later that for each d, q, r, and l (where r, the length of our sentence, is relevant since $k \leq r$, and where l is the number of built-in relations) the number of d, q-colors is finite, and an upper bound depends only on d, q, l, and r (and not, for example, on n, the number of vertices). Note also that the d, q-color of a vertex determines its d, q'-color if $q' < q$. The rough idea as to why we are interested in d, q-colorings is as follows: Assume that the spoiler and duplicator are playing an r-round Ehrenfeucht-Fraïssé game on G_0 and G_1. Assume that through the first j rounds, the points $a_1, \ldots, a_j$ picked in G_0 have the same d, $q' + 1$-colors as the corresponding points $b_1, \ldots, b_j$ picked in G_1, where $q' \geq 1$ and $d > j$. If the spoiler then picks, say, a_{j+1} in G_0 where a_{j+1} is adjacent to a_i in G_0, for some $i \leq j$, then the duplicator is guaranteed that there is a point b_{j+1} in G_1 that is adjacent to b_i in G_1, and where the d, q'-color of b_{j+1} is the same as the d, q'-color of a_{j+1}. If a_{j+1} is distinct from any of $a_1, \ldots, a_j$, and if $d > j$, then b_{j+1} can be taken to be distinct from any of $b_1, \ldots, b_j$.

Our goal in the remainder of this paper is to prove Proposition 8.2, which says that the duplicator has a winning strategy in the r-round Ehrenfeucht-Fraïssé game played on G_0 and G_1, where the isomorphisms must respect vertex and edge colors. We now give a sketch of how we shall prove this result.

Three important parameters that we shall select are ξ, σ, and θ, where ξ and θ will be selected to be large, and σ to be small (but positive). The graph Γ contains no cycle of length less than ξ, and the degree of each point in Γ is at most $|V|^\sigma$. The probability p of a backedge lies in the range $\theta/n < p < n^\sigma/n$. We shall show that if θ is sufficiently large (so that the probability p of a backedge is sufficiently large), then, with high probability, each vertex has the same d, q-color (for appropriate d, q) in G_0 as in G_1. The reason this is true is that if $e = \langle x, y \rangle$ is the deleted forward edge, and if p is large enough, then there are enough backedges from x to "play the role" of the missing edge e. One reason it is useful that (with high probability) each vertex has the same d, q-color in G_0 as in G_1 is that in the first round, whatever point the spoiler selects, the duplicator can select the same point in the opposite graph, and be guaranteed that (with high probability), these points have the same d, q-color.

Let $m = \lceil \xi/2 \rceil$. We shall show that if σ is sufficiently small, then, with high probability, each neighborhood of radius m in G_0 is small. This is because the degree of each point in Γ is at most $|V|^\sigma$, and because the probability p of a backedge is made small enough. As we shall show, it follows that, with high probability, the number of points in short cycles (when we ignore the directions of edges) is small.

Let A be the set of special points (those points on short cycles, and those points whose color is unusual, in a sense to be specified later). From what we have said, we can select our parameters so that, with high probability, A is small. Let C be the set of points near (within distance m) of A. Since A is small, and since neighborhoods of radius m are small, it follows that C is small. In particular, with high probability, neither endpoint of the deleted forward edge is in C. So we assume that neither endpoint of the deleted forward edge is in C.

The duplicator's winning strategy is roughly as follows. When the spoiler picks a point u, the duplicator picks the same point u in the other graph, unless u is near some point in $\bar{C}$ (the complement of C) that has already been picked by one of the players. What if the spoiler picks a point that is near some point in $\bar{C}$ that has already been picked by one of the players? For definiteness, let us say that on round 35, the spoiler selects the point a_{35} in G_0, that is distance 3 from the point a_{17} that was picked earlier. Say that v_0, v_1, v_2, v_3 is a minimal path in $U(G_0)$ from a_{17} to a_{35}, where $v_0 = a_{17}$ and $v_3 = a_{35}$. We assume by inductive assumption that a_{17} has the same d, q-coloring in G_0 as the point b_{17} has in G_1, for some large d and q. Assume for definiteness that, say, there is a directed edge $\langle v_1, v_0 \rangle$ but no directed edge $\langle v_0, v_1 \rangle$ in

G_0. Let the $d, q-1$-color of v_1 be Ω, and let the C_0-color of the edge $\langle v_1, v_0\rangle$ be Ω'. Since a_{17} and b_{17} have the same d, q-color, we have that the number of points in G_0 with $d, q-1$-color Ω and that have an edge in G_0 with C_0-color Ω' to a_{17} but no reverse edge is either greater than d, or the same as the number of points in G_1 with $d, q-1$-color Ω and that have an edge in G_1 with C_0-color Ω' to b_{17} and no reverse edge. Thus, if d is sufficiently large, then there is a point w_1 in G_1, distinct from any previously selected b_i, with $d, q-1$-color Ω and that has an edge in G_1 with C_0-color Ω' to b_{17}, and no reverse edge. Similarly, we find w_2 and w_3, and let $b_{35} = w_3$. Then a_{35} and b_{35} have the same $d, q-3$-color. Since there are no small cycles in $\bar{C}$, the subgraph of G_0 induced by v_0, v_1, v_2, v_3 is isomorphic to the subgraph of G_1 induced by w_0, w_1, w_2, w_3, and the isomorphism respects color.

Unfortunately, in order to give a precise proof along the lines of the ideas in the sketch we just gave, we have to do considerable work, which we now begin. We note that the winning strategy for the duplicator that we shall give is more complicated than the strategy we just sketched.

In the next lemma, the order of choice of Q_p^n and $A_1, \ldots, A_k$ is reversed. Thus, we can think of the situation in Proposition 8.2 as one where the sets $A_1, \ldots, A_k$ are selected by an adversary after Q_p^n has been randomly generated. However, in Lemma 8.3 below, the adversary must select $A_1, \ldots, A_k$ before seeing what Q_p^n is. This is an important difference. For example, if we know that, say, $u \in A_1$ and $v \notin A_1$, then in the latter case (the situation of Lemma 8.3), the probability that there is a backedge with endpoints u and v is simply p (the probability of an arbitrary backedge). However, in the former case, this is not necessarily true: it is possible, for example, that A_1 has been selected so that there are no backedges with endpoints in A_1.

LEMMA 8.3. *Assume that d, q, r, l, and ξ are positive integers, and $\varepsilon > 0$. Assume that θ is sufficiently large with respect to d, q, r, l, and ε, that $0 < \sigma < 1/2$, and that n is sufficiently large with respect to r and ε. Suppose that V is a set with n elements, and that $\langle B_1, \ldots, B_l\rangle$ is a successor-like family with parameters ξ and σ. Let $A_1, \ldots, A_k$ be arbitrary unary relations over V, where $k \le r$. Assume that $\theta/n < p < n^\sigma/n$, and that Q_p^n is a random path on V with random backedges. Then with probability at least $1 - \varepsilon 2^{-kn}$ (where the probability is taken over the choices of Q_p^n), the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that each vertex has the same d, q-color in G_0 as in G_1.*

PROOF. Let us define the *d, q-edge color* of a pair $\langle v, w\rangle$, where $v, w \in V(G_0)$, to be a complete description of the edge color of $\langle v, w\rangle$ and the d, q-colors of the vertices v and w. As before, there is a finite upper bound η (that depends only on d, q, l, and r) on the number of d, q-edge colors. Let τ be a positive real number chosen sufficiently small with respect to η and ε, and let θ be sufficiently large with respect to d and τ. Let us say that a pair $\langle v, w\rangle$ is *Γ-free* if (v, w) is not an edge of Γ. Call a forward edge f of G_0 *special* if (a) f is not Γ-free, or if (b) when we denote the d, q-edge color of f by Ω, then f is among the first $\lceil \tau n\rceil$ or last $\lceil \tau n\rceil$ of the forward edges of G_0 with d, q-edge color Ω.

For each point v, the probability that the forward edge f starting at v satisfies (a) (i.e., is not Γ-free) is at most n^σ/n, since the degree of v in Γ is at most n^σ. If $\sigma < 1/2$, we can make this probability be less than $\varepsilon/4$ for n sufficiently large with respect to ε. Also, the number of forward edges f that satisfy (b) is at most $2\lceil \tau n\rceil \eta$, so, by taking τ

sufficiently small with respect to η and ε, we can make the probability that a randomly selected forward edge f satisfy (b) be less than $\varepsilon/4$. So by suitable choice of parameters, the probability that a randomly selected forward edge f is special is less than $\varepsilon/2$.

Let E be an arbitrary set of $\lceil \tau n \rceil$ vertices, and let u be an arbitrary vertex not in E. Let us say that u is *weak for* E if it is not the case that there are at least d Γ-free backedges b where one endpoint of b is u and the other endpoint v is a member of E (as before, we say that such a backedge *involves* u and v). The probability of a Γ-free backedge involving u and a given member of E is at least $p(1 - n^{\sigma}/n)$, since the degree of u in Γ is at most n^{σ}. If $\sigma < 1/2$ and $n \geq 4$, then this probability is at least $p/2$. So the expected number of Γ-free backedges involving u and a point in E is at least $p\tau n/2 > \theta\tau/2$ (since $p > \theta/n$). By taking θ sufficiently large with respect to d and τ, this expected number is at least $2d$. By Corollary 7.2 to the large deviation theorem, if θ is sufficiently large that $\theta\tau/2 > c_1$, then the probability that there are less than d backedges involving u and a point in E is less than $e^{-c_2\theta\tau/2}$. By taking θ sufficiently large with respect to τ, we can make this probability as small as we want. Let us call this probability δ (shortly, we shall select how small we want δ to be). If A is a set of at least $n\varepsilon/(4\eta)$ vertices, and E is a set of $\lceil \tau n \rceil$ vertices, and if A and E are disjoint, then let $Y(A, E)$ be the event that every member of A is weak for E. So the probability of $Y(A, E)$ is at most $\delta^{n\varepsilon/(4\eta)}$. Since there are at most 2^n possibilities for A and 2^n possibilities for E, the probability that there exists some pair A, E such that $Y(A, E)$ occurs is at most

$$2^n 2^n \delta^{n\varepsilon/(4\eta)} = 2^{n(2 + (\log_2 \delta)(\varepsilon/(4\eta)))}.$$

We can take δ sufficiently small with respect to ε and η that $2 + (\log_2 \delta)(\varepsilon/(4\eta)) < -(k + 1)$, so that this probability is less than $2^{-(k+1)n}$. We can then assume that n is sufficiently large with respect to r (and hence with respect to $k \leq r$) that this probability is less than $\varepsilon 2^{-kn}$.

So with probability at least $1 - \varepsilon 2^{-kn}$, it is the case that Q_p^n has been selected so that $Y(A, E)$ fails for every A, E. So assume now that $Y(A, E)$ fails for every A, E. That is, for every set E of $\lceil \tau n \rceil$ vertices, there are less than $n\varepsilon/(4\eta)$ vertices that are weak for E.

Let us say that a forward edge $f = \langle u, v \rangle$ is *bad* if it is not the case that (a) there are at least d Γ-free backedges from u to points with the same d, q-color as v, and (b) there are at least d Γ-free backedges to v from points with the same d, q-color as u. We now count the number of nonspecial edges f that are bad. Assume that f is not special. Let Ω be the d, q-edge color of f. Since f is not special, f is not among the first or last $\lceil \tau n \rceil$ forward edges of G_0 with d, q-edge color Ω. Let E_1 be the set of the first $\lceil \tau n \rceil$ forward edges of G_0 with d, q-edge color Ω, and let E_2 be the set of the last $\lceil \tau n \rceil$ forward edges of G_0 with d, q-edge color Ω. Then f comes after the edges of E_1 (that is, if z is the head of an arbitrary edge of E_1, then $x > z$), and f comes before the edges of E_2. Now there are less than $n\varepsilon/(4\eta)$ vertices that are weak for E_1, and less than $n\varepsilon/(4\eta)$ vertices that are weak for E_2. So the number of forward edges $f = \langle u, v \rangle$ where u is weak for E_1 or v is weak for E_2 is less than $n\varepsilon/(2\eta)$. Therefore, the number of nonspecial forward edges with d, q-edge color Ω that are bad is less than

$n\varepsilon/(2\eta)$. So the total number of nonspecial forward edges that are bad is less than $n\varepsilon/2$.

From what we have shown, it follows that the probability is at least $1 - \varepsilon 2^{-kn}$ (where the probability is taken over the choices of Q_p^n) that the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that e is neither special nor bad. If $e = \langle x, y \rangle$ is neither special nor bad, then (a) e is Γ-free, (b) there are at least d Γ-free backedges $\langle x, w \rangle$ where the d, q-color of w is the same as the d, q-color of y, and (c) there are at least d Γ-free backedges $\langle w, y \rangle$ where the d, q-color of w is the same as the d, q-color of x. So assume that (a), (b), and (c) hold. We show by induction on $q' \leq q$ that the vertices of G_0 and G_1 have the same d, q'-color. For $q' = 1$ this is by definition, since the d, 1-coloring on G_0 (resp. G_1) is simply the C_0 (resp. C_1) color, and, as we have noted, these colorings are the same on vertices. Suppose inductively that the d, q'-colorings are identical for some q', where $1 \leq q' < q$. The $d, q' + 1$-color of a point z is determined by the d, q'-color of z along with the number of edges of each C_0-color (resp. C_1-color) going to and the number of such edges coming from points with a fixed d, q'-color. Assume now that $z \neq x$ and $z \neq y$ (where $e = \langle x, y \rangle$ is the deleted forward edge). By our induction assumption, the d, q'-color of z and its neighbors are identical in both G_0 and G_1; since also the C_0 and C_1 colors of edges containing z are identical, it follows that the $d, q' + 1$-colors of z are also identical.

We now consider the $d, q' + 1$-color of x. The only difference between the graphs G_0 and G_1 is the edge $e = \langle x, y \rangle$. However, e is Γ-free, and by assumption there are at least d Γ-free backedges $\langle x, w \rangle$ where the d, q-color of w is the same as the d, q-color of y (and hence the d, q'-color of w is the same as the d, q'-color of y). So all of these edges will be counted together with $\langle x, y \rangle$. Since their number is at least d, the omission of e will not change the $d, q' + 1$-color of x. (Here we also used the fact that the colorings do not distinguish between forward edges and backedges.) So x has the same $d, q' + 1$-color in both G_0 and G_1. A similar argument shows that y also has the same $d, q' + 1$-color in both G_0 and G_1. This completes the induction step. ∎

We need Lemma 8.3 only to prove Lemma 8.4 below. Unlike Lemma 8.3, in Lemma 8.4 the unary relations $A_1, \ldots, A_k$ can be thought of as being selected by an adversary after Q_p^n has been randomly generated (as in Proposition 8.2). Lemma 8.3 says that the probability that a *fixed* sequence $A_1, \ldots, A_k$ is "good" is high, and Lemma 8.4 says that the probability that *every* sequence $A_1, \ldots, A_k$ is "good" is high (see the proof of Lemma 8.4 to see what "good" means). In Lemma 8.4 and in all of the other remaining lemmas, we assume that V is a set with n elements, that $\langle B_1, \ldots, B_l \rangle$ is a successor-like family with parameters ξ and σ, that $\theta/n < p < n^{\sigma}/n$, that Q_p^n is a random path on V with random backedges, and that $A_1, \ldots, A_k$ are arbitrary unary relations over V (where $k \leq r$), selected after Q_p^n. Lemma 8.4 makes good our claim that, with high probability, each vertex has the same d, q-color in G_0 as in G_1.

LEMMA 8.4. *Assume d, q, r, l, and ξ are positive integers, and $\varepsilon > 0$. Assume that θ is sufficiently large with respect to d, q, r, l, and ε, that $0 < \sigma < 1/2$, and that n is sufficiently large with respect to r and ε. Then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), for each choice of the unary relations $A_1, \ldots, A_k$ over V, the probability is at least $1 - \varepsilon$ (where the probability is taken over* *the choices of the deleted forward edge e) that each vertex has the same d, q-color in G_0 as in G_1.*

PROOF. Let us say that $A_1, \ldots, A_k$ is *good for* $\mathscr{M} = \langle V, Q_p^n, s, t, B_1, \ldots, B_l \rangle$ if the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that each vertex has the same d, q-color in G_0 as in G_1. By Lemma 8.3, we know that the probability is at most $\varepsilon 2^{-kn}$ that $A_1, \ldots, A_k$ is not good for M. Since there are only 2^{kn} possible choices for $A_1, \ldots, A_k$, it follows that the probability is at most $\varepsilon 2^{-kn} 2^{kn} = \varepsilon$ that there is some $A_1, \ldots, A_k$ that is not good for M. So with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), every $A_1, \ldots, A_k$ over V is good for M. ■

If G is a (possibly colored) directed graph, let $U(G)$ be the ordinary undirected graph which is obtained by ignoring the colors of vertices and edges, and ignoring the directions of edges. Whenever we speak of the distance $d_G(v, w)$ between two points v and w in G, we mean the distance in $U(G)$. We may refer to this distance as the *G-distance*.

The next lemma makes good our claim that neighborhoods of radius $m = \lceil \xi/2 \rceil$ are small.

LEMMA 8.5. *Let ξ be a positive integer, let $m = \lceil \xi/2 \rceil$, and let $\varepsilon > 0$. If σ is sufficiently small with respect to ξ, and n is sufficiently large with respect to ε and σ, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), for every vertex $v \in V(G_0)$, the set of points whose G_0-distance from v is at most m is less than $n^{1/100}$.*

PROOF. Let v be a fixed vertex. If $0 < q < 1$, then define P_q to be the probability that there are at least $2n^\sigma$ backedges to or from v, if the probability of each backedge appearing were q (thus, we are actually interested in P_p, where p is our probability of selecting each backedge). It is clear that P_q is monotone in q, so that $P_p \leq P_{n^\sigma/n}$ (since $p < n^\sigma/n$). If the probability of each backedge appearing were n^σ/n, then the expected number of backedges would be n^σ. So if n is sufficiently large with respect to σ, then, by Corollary 7.2 to the large deviation theorem, $P_{n^\sigma/n} < e^{-c_2 n^\sigma}$. Hence, $P_p < e^{-c_2 n^\sigma}$. So, with probability at most $ne^{-c_2 n^\sigma}$, for some point v the number of backedges to or from v is at least $2n^\sigma$. If n is sufficiently large with respect to ε and σ, then this probability is less than ε. So assume that the number of backedges to or from each point is less than $2n^\sigma$. Then the degree of each point in G_0 is less than $3n^\sigma + 2$ (where the extra n^σ is from Γ, and the 2 is from the forward edges). This is less than $n^{2\sigma}$ for n sufficiently large with respect to σ. So the number of points in a neighborhood of radius m is less than $1 + n^{2\sigma} + n^{4\sigma} + \cdots + n^{2m\sigma}$, which is a geometric series bounded above by $n^{2(m+1)\sigma}$. If σ is sufficiently small with respect to m, then this is less than $n^{1/100}$. ■

The next lemma makes good our claim that the number of points in short cycles (when we ignore the directions of edges) is small.

LEMMA 8.6. *Let ξ be a positive integer, let $m = \lceil \xi/2 \rceil$, and let $\varepsilon > 0$. If σ is sufficiently small with respect to ξ, and if n is sufficiently large with respect to ε and σ, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), the number of points that are on some cycle shorter than m in the graph $U(G_0)$ is less than $n^{3/4}$.*

PROOF. By Lemma 8.5, if σ is sufficiently small with respect to ξ, and n is sufficiently large with respect to ε and σ, then, with probability at least $1 - \varepsilon/8$, every neighborhood of radius m contains less than $n^{1/100}$ points. But if it were the case that every neighborhood of radius m contains less than $n^{1/100}$ points, then the degree of each vertex would be less than $n^{1/100}$, so the number of edges would be less than $n^{1+1/100}$. Hence, with probability at least $1-\varepsilon/8$, the number of edges would be less than $n^{1+1/100}$. Let us refer to a cycle of length shorter than m in $U(G_0)$ as a *short cycle*. Let g denote the number of backedges in Q_p^n, and W the set of points v such that v is on some short cycle. Since the probability is at least $1-\varepsilon/8$ that the number of edges is less than $n^{1+1/100}$, it follows easily that we need only show that the conditional probability that $|W|<n^{3/4}$, given that the number of edges is less than $n^{1+1/100}$, is at least $1-\varepsilon/2$. So we need only show that the conditional probability that $|W|<n^{3/4}$, given that $g=c$, is at least $1-\varepsilon/2$ for all $c=1,\ldots,n^{1+1/100}$. Suppose now that $g=c$, for some fixed $c=1,\ldots,n^{1+1/100}$. This means that Q_p^n contains exactly g backedges with uniform distribution. We may get these g backedges by picking a sequence $f_1,\ldots,f_g$ of backedges from the edges of the complete graph, so that if $E_i=\{f_1,\ldots,f_i\}$ for $1\le i\le g$, then the distribution of f_i is uniform on the complement of E_{i-1}.

Let us refer to a short cycle in the graph consisting of Γ and the undirected version of all of the backedges as a *type* 1 *short cycle*. Let us say that the backedge f_i *causes a type* 1 *short cycle* if the undirected version of f_i is an edge of a type 1 short cycle in $\Gamma\cup U(E_i)$. Further, let us say that the backedge f_i is *bad* if (a) there is no point whose neighborhood of radius m in $\Gamma\cup U(E_{i-1})$ is of size at least $n^{1/100}$, and (b) f_i causes a type 1 short cycle.[4] We now give an upper bound on the probability that f_i is bad for a fixed i with $1\le i\le g$. If there is some point whose neighborhood of radius m in $\Gamma\cup U(E_{i-1})$ is of size at least $n^{1/100}$, then the probability that f_i is bad is 0. So assume that the neighborhood of radius m in $\Gamma\cup U(E_{i-1})$ of every point is of size less than $n^{1/100}$. Then the number of pairs of points that are of distance less than m apart in $\Gamma\cup U(E_{i-1})$ is at most $n^{1+1/100}$, since each of the n points has at most $n^{1/100}$ points within distance m. However, the total number of pairs of points that are not in E_{i-1} is at least $C_2^n-n^{1+1/100}$, where C_2^n is the binomial coefficient, and where $n^{1+1/100}$ is an upper bound on the size of E_{i-1} (since the number of edges in $\Gamma\cup U(E_{i-1})$ is at most $n^{1+1/100}$). Therefore the probability that f_i is bad in this case is at most $n^{1+1/100}/(C_2^n-n^{1+1/100})$, which is less than $n^{-1/2}$ if n is sufficiently large. So the probability that f_i is bad is less than $n^{-1/2}$, for each fixed $i=1,\ldots,g$. Hence the expected number of bad backedges is less than $n^{1+1/100}n^{-1/2}<n^{2/3}$. So by Corollary 7.2 to the large deviation theorem, if n is sufficiently large with respect to ε, then with probability at least $1-\varepsilon/8$ the number of bad backedges is less than $2n^{2/3}$. So assume for now that the number of bad backedges is less than $2n^{2/3}$. But also, as we noted, with probability at least $1-\varepsilon/8$, every neighborhood of radius m contains less than $n^{1/100}$ points. So assume for now that every neighborhood of radius m contains less than $n^{1/100}$ points. It follows that if f_i causes a type 1 short cycle, then f_i is a bad backedge. So it follows that the number of backedges that cause a type 1

[4]The reader may wonder why we do not simply invoke Lemma 8.5 and say that, with high probability, every neighborhood of radius m in G_0 contains less than $n^{1/100}$ points, assume that this is true, and go on from there. The reason is that if we were to make this assumption, then it would no longer be the case that the distribution of f_i is uniform on the complement of E_{i-1}, since this assumption has a subtle influence on the probabilities.

short cycle is less than $2n^{2/3}$. Now every point in a type 1 short cycle caused by f_i is of distance at most m from the tail of f_i. But by assumption, there are less than $n^{1/100}$ such points. So the total number of points on type 1 short cycles is at most $2n^{2/3}n^{1/100}$, which is less than $n^{3/4}/2$ if n is sufficiently large.

Similarly, we can define a *type* 2 *short cycle* to be a short cycle that contains a forward edge. Let $f_1, \ldots, f_{n-1}$ now be the forward edges in order (thus, the head of f_i is the tail of f_{i+1} for $1 \leq i < n - 1$). Let $F_i = \{f_1, \ldots, f_i\}$ for $1 \leq i < n$, and let E be the set of backedges. We say that a forward edge f_i *causes a type* 2 *short cycle* if the undirected version of f_i is an edge of a type 2 short cycle in $\Gamma \cup U(E) \cup U(F_i)$. We can now consider the forward edges one by one in order, and estimate the probability that the forward edge causes a type 2 short cycle. Here, the argument is a little easier. By a very similar argument to that contained in the proof of Lemma 8.5, the probability is at least $1 - \varepsilon/8$ that every point in $\Gamma \cup U(E)$ has degree at most $n^{1/(101(m+1))}$. So assume for now that every point in $\Gamma \cup U(E)$ has degree at most $n^{1/(101(m+1))}$. Then the degree of every point in G_0 is at most $2 + n^{1/(101(m+1))}$, which is less than $n^{1/(100(m+1))}$ if n is sufficiently large with respect to m. So by summing a geometric series as in the proof of Lemma 8.5, every neighborhood of radius m in G_0 contains less than $n^{1/100}$ points.[5] We now consider which of the first $(n - 1) - n^{1/2}$ forward edges cause a type 2 short cycle. The probability that such a forward edge causes a type 2 short cycle is less than $n^{1/100}/n^{1/2}$, so, by our usual argument, with probability at least $1 - \varepsilon/8$, the number of such edges is less than $2(n - 1)n^{1/100}/n^{1/2} < 2n^{2/3}$. So even if all of the last $n^{1/2}$ forward edges cause a type 2 short cycle, the total number of forward edges that cause a type 2 short cycle is less than $2n^{2/3} + n^{1/2}$. So, as before, the total number of points on type 2 short cycles is at most $(2n^{2/3} + n^{1/2})n^{1/100}$, which is less than $n^{3/4}/2$ if n is sufficiently large. Since a point is on a short cycle iff it is on a type 1 short cycle or a type 2 short cycle, our result follows. ■

We now define some more concepts we will need. Let G be a colored directed graph, D a subset of $V(G)$, and j a positive integer. We say that the vertex color Ω is *j-dispersed in D* if there is a subset R of D consisting entirely of points of color Ω, so that $|R| = j$ and for all $v, w \in R$ we have $d_G(v, w) > j$.

We say that the partition $T = \langle A, B, W\rangle$ of the set $V(G)$ into three subsets is an *m-trisection* with respect to the coloring, where m is a positive integer, if the following conditions are satisfied:

1. Each point in each cycle shorter than m in the graph $U(G)$ is contained in A.
2. If Ω is a color that is not m-dispersed in W, then all of the points of color Ω are contained in A.
3. If $a \in A$ and $w \in W$, then $d_G(a, w) > m$.

Intuitively, in an m-trisection $\langle A, B, W\rangle$, the set A contains all of the special points

[5] As with the backedges, we cannot simply invoke Lemma 8.5 and say that, with high probability, every neighborhood of radius m in G_0 contains less than $n^{1/100}$ points, assume that this is true, and go on from there. Similarly to before, if we were to make this assumption, then it would no longer be the case that the selection of f_{i+1} is made independently of $\Gamma \cup U(E) \cup U(F_i)$. However, assumptions on degrees in $\Gamma \cup U(E)$ do not effect probabilities of forward edges, since the forward edges are selected independently of $\Gamma \cup U(E)$.

(those on short cycles, and those whose color is unusual), the set W contains typical points, and B is the borderline between A and W.

The proof of the next lemma shows that the set A of special points is small (as promised earlier). We do not bother to state explicitly in the statement of the lemma that A is small, but instead state only those facts that we need later.

LEMMA 8.7. *Let d, q, r, and l be positive integers, and $\varepsilon > 0$. If ξ is sufficiently large with respect to ε, if $m = \lceil \xi/2 \rceil$, if σ is sufficiently small with respect to ξ, and n is sufficiently large with respect to $r, l, \varepsilon, \xi, \sigma, d$, and q, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), there is an m-trisection $\langle A, B, W \rangle$ of G_0 with respect to the d, q-coloring such that if X is the set of points in $V(G_0)$ whose G_0-distance from $A \cup B$ is at most m, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e), X does not contain the tail or the head of the deleted forward edge e.*

PROOF. We will define a sequence of increasing sets Y_j, by induction on j. The last element of the sequence will be A.

First, we define Y_1 to be the set of points that are on some cycle shorter than m in the graph $U(G_0)$. By Lemma 8.6, we can assume that $|Y_1| < n^{3/4}$. Since we will define A with $A \supseteq Y_1$, the definition of Y_1 already implies that A will satisfy the first condition in the definition of an m-trisection for G_0. Suppose that $Y_1 \subseteq \cdots \subseteq Y_j$ is already defined. Let N_j be the set of those points whose G_0-distance from Y_j is at most m. Let X_j be the set of those points of $V(G_0) - Y_j$ whose d, q-color is not m-dispersed in the set $V(G_0) - N_j$. If X_j is not empty, then let $Y_{j+1} = Y_j \cup X_j$. It is easy to see that the number of colors outside Y_j strictly decreases with j; this is because Y_{j+1} is defined by taking all points of some color Ω that is represented outside of Y_j, and putting all points with color Ω into Y_{j+1}. So for some z, which is at most one plus the number of d, q-colors, $X_z = \varnothing$ (recall that the number of d, q-colors depends only on d, q, r and l, and not on n or σ). Let $A = Y_z$, $B = N_z - A$, and $W = V(G_0) - N_z$. It is clear that our construction produces an m-trisection.

We now show that if X is as in the statement of the lemma, then $|X| < n^{4/5}$ for sufficiently large n. To do so, we first estimate the size of A. By Lemma 8.5, we can assume that each neighborhood of radius m contains less than $n^{1/100}$ points. Let us estimate $|X_j| = |Y_{j+1} - Y_j|$, where $1 \leq j < z$. Recall that X_j is the set of those points of $V(G_0) - Y_j$ whose d, q-color is not m-dispersed in the set $V(G_0) - N_j$. Assume that Ω is a d, q-color that is not m-dispersed in the set $V(G_0) - N_j$. Let $X_j(\Omega)$ be those points of color Ω in X_j. Let Z be a maximal subset of $X_j(\Omega)$ with the property that every pair of distinct points in Z are of distance greater than m apart (if every pair of distinct points in $X_j(\Omega)$ is within distance m, then Z is a singleton set). Since Ω is not m-dispersed in $V(G_0) - N_j$, we know that $|Z| < m$. By maximality, $X_j(\Omega)$ is covered by the union of the neighborhoods of radius m of each point in Z. Each such neighborhood contains less than $n^{1/100}$ points; hence, since there are less than m such neighborhoods, $|X_j(\Omega)| < mn^{1/100}$. So $|X_j|$ is bounded by the number η of d, q-colors times this bound on $X_j(\Omega)$. Thus, $|X_j| < \eta m n^{1/100}$. So $|A - Y_1| < z\eta m n^{1/100} \leq (\eta + 1)\eta m n^{1/100}$. Since $|Y_1| < n^{3/4}$, it follows that $|A| < n^{4/5}$ if n is sufficiently large with respect to r, ξ, d, q, and l (and hence with respect to m and η).

Now B is defined to contain those points not in A whose G_0-distance from A is at most m. So B is the union of the neighborhoods of radius m about the points in A.

Each such neighborhood contains less than $n^{1/100}$ points, and there are $|A| < n^{4/5}$ such neighborhoods, which tells us that $|A \cup B| < n^{4/5}n^{1/100}$. Similarly, $|X| < |A \cup B|n^{1/100} < n^{4/5}n^{1/100}n^{1/100} < n^{5/6}$. We have shown that with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), $|X| < n^{5/6}$.

Let $e = \langle x, y \rangle$ be the deleted forward edge. Since there are $n - 1$ forward edges, and since $|X| < n^{5/6}$, the probability that $x \in X$ is less than $n^{5/6}/n$, and similarly for y. So if n is sufficiently large with respect to ε, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e), the set X does not contain the tail or the head of the deleted forward edge e. ■

Before we describe the duplicator's winning strategy (and thereby prove Proposition 8.2) we need a few more definitions. First, it is convenient to slightly modify the definition of the rules of the game to an equivalent but technically more convenient form.

Let G_0 and G_1 be colored graphs. Let r be a positive integer. We now define an *r-round Ehrenfeucht-Fraïssé game (played on G_0 and G_1)*. There are two players (the spoiler and the duplicator), and two graphs, G_0 and G_1. On round j, the spoiler selects a point x_j from the universe of one of G_0 or G_1. (Since the universes of G_0 and G_1 need not be disjoint, the spoiler not only selects a point but labels it with either G_0 or G_1. If the spoiler selects the point x_j and labels it with, say, G_0, then we say that *the spoiler selected x_j from G_0*.) The duplicator then selects a partial isomorphism λ_j from a subset of the universe of G_0 to a subset of the universe of G_1. By *partial isomorphism*, we mean (a) λ_j is a one-to-one map; (b) if v is in the domain of λ_j, then the colors of v and $\lambda_j(v)$ are the same, and (c) if v and w are distinct points in the domain of λ_j, then there is an edge $\langle v, w \rangle$ of color Ω in G_0 iff there is an edge $\langle \lambda_j(v), \lambda_j(w) \rangle$ of color Ω in G_1. There are two restrictions on λ_j:

1. If $j > 1$, then λ_j is an extension of the partial isomorphism λ_{j-1} selected in the previous round.
2. If the spoiler selected x_j from G_0 (resp. G_1), then the domain (resp. range) of λ_j contains x_j.

The duplicator wins if each λ_j that he selects (for $1 \leq j \leq r$) satisfies these restrictions; otherwise, the spoiler wins.

Assume that α and β are positive integers (where, for convenience, we assume that $\beta \geq 2$). Assume that $H \subseteq V(G)$ (in the cases of interest to us, G is either G_0 or G_1), and that P is a partition of H. We say that P is an *α, β-decomposition* of H if the following conditions are satisfied:

1. Each class of P is the set of points in a connected component of the subgraph of $U(G)$ generated by H (so, in particular, each case of P is a connected subset of the graph $U(G)$).
2. If v and w are in the same class of P, then their G-distance is less than α.
3. If v and w are in different classes of P, then their G-distance is at least β.

Let λ be a partial isomorphism from G_0 to G_1, where the domain of λ is H. We say that λ is *α, β-perfect* if the following conditions hold, when P is the partition of H where each class is the set of points in a connected component of the subgraph of $U(G)$ generated by H. We assume that $\langle A, B, W \rangle$ is a fixed m-trisection of G_0 (so "α, β-perfect" is really defined with respect to $\langle A, B, W \rangle$).

1. P is an α, β-decomposition of H in G_0.

2. $\lambda(P)$ is an α, β-decomposition of $\lambda(H)$ in G_1 (by $\lambda(P)$, we mean the obvious partition of the range of λ).

3. If ρ is a class of P so that either ρ or $\lambda(\rho)$ is contained in $A \cup B$, then the restriction of λ onto ρ is the identity map.

We are now ready to begin the proof of Proposition 8.2. We must show that with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), for each choice of the unary relations $A_1, \ldots, A_k$ over V, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that the duplicator has a winning strategy in the r-round Ehrenfeucht-Fraïssé game played on G_0 and G_1, where the isomorphisms must respect the vertex and edge colorings as given by C_0 and C_1 (we use the more refined d, q-coloring on vertices only as a tool in the proof). For convenience, we introduce a 0th round, where no points have been selected so far by either player, and so the domain of the partial isomorphism λ_0 after round 0 is the empty set. Let f be the function $x \mapsto 2^x$, with domain the positive integers. We define f^j inductively by letting f^1 be f, and $f^{j+1} = f \circ f^j$. Let $d = f^{3^{r+3}}(1) + 1$. Assume that ξ (and hence $m = \lceil \xi/2 \rceil$) is sufficiently large that $m > d$. Define $q_j = (r - j)(f^{3^{r+3}}(1) + 1) + 1$, for $0 \le j \le r$, and $q = q_0$. Thus, d and q depend only on r. By Lemmas 8.4 and 8.7, we know that if $\varepsilon > 0$, if ξ is sufficiently large with respect to ε, if θ is sufficiently large with respect to r, l, and ε, if $\sigma > 0$ is sufficiently small with respect to ξ, and if n is sufficiently large with respect to r, l, ε, ξ and σ, then with probability at least $1 - \varepsilon$ (where the probability is taken over the choices of Q_p^n), for each choice of the unary relations $A_1, \ldots, A_k$ over V, where $k \le r$, the probability is at least $1 - \varepsilon$ (where the probability is taken over the choices of the deleted forward edge e) that (a) each vertex has the same d, q-color in G_0 as in G_1, and (b) there is an m-trisection $\langle A, B, W \rangle$ of G_0 with respect to the d, q-coloring such that if X is the set of points in $V(G_0)$ whose G_0-distance from $A \cup B$ is at most m, then X does not contain the tail or the head of deleted forward edge e. So assume that (a) and (b) hold. Define a *partial d, q′-isomorphism* to be a partial isomorphism that preserves d, q'-colors. We prove inductively on j that, by proper play, the duplicator can maintain the following conditions after round j (for $0 \le j \le r$), where we take $\alpha_{-1} = 0$ and $\beta_{-1} = \infty$.

1. λ_j is a partial d, q_j-isomorphism whose domain (resp. range) contains the point selected by the spoiler if the spoiler's move in round j is to select a point in G_0 (resp. G_1), and if $j > 0$ then λ_j extends λ_{j-1}.

2. There are positive integers α_j and β_j such that
   (a) $\alpha_{j-1} < \alpha_j < f^{3^{r-j+3}}(\alpha_j) < \beta_j < \beta_{j-1}$,
   (b) the domain H_j of λ_j is of size at most α_j, and
   (c) λ_j is α_j, β_j-perfect.

It is clear that condition (1) is sufficient to guarantee that the duplicator has a winning strategy, since each partial d, q_j-isomorphism is also a partial d, 1-isomorphism.

For the base case $j = 0$, we let λ_0 have empty domain, $\alpha_0 = 1$, and $\beta_0 = f^{3^{r+3}}(1) + 1 < m$ (we can think of the partition of the domain as having only one class, the empty set). Assuming that the conditions hold after round j, we now show that the duplicator can move in such a way that the conditions hold after round $j + 1$ (for $0 \le j < r$). In our arguments below, we shall sometimes speak only of the "distance" between points, when it is clear whether we mean the G_0-distance or the G_1-distance, or when it does not matter. We shall sometimes take advantage of the following facts:

1. If $a \in A \cup B$, then $d_{G_0}(a, b) = \gamma < m$ iff $d_{G_1}(a, b) = \gamma < m$ (this is because neither the head nor the tail of the deleted forward edge e is in X).

2. The m-trisection $\langle A, B, W \rangle$ of G_0 is also an m-trisection of G_1 (this follows from the fact that the G_0-distance between any pair of points is less than or equal to the G_1-distance).

Assume that on round $j + 1$, the spoiler selects point v in G_0 (by the two facts just given and by the symmetry of our induction assumption, the reader can verify by going through the proof below that a very similar argument holds if the spoiler selects a point in G_1). Let κ be the G_0-distance of v and H_j, and let ρ be a class in the partition of H_j that is of G_0-distance κ from v. Define i to be 3^{r-j+3}, and let $\kappa_0 = f^{\lfloor i/2 \rfloor}(\alpha_j)$.

There are five cases. The first two cases (Cases I and II) occur when the point v selected by the spoiler is near (within distance κ_0 of) some point in G_0 that was selected earlier. In Case I, the nearest class ρ to v is completely contained in $A \cup B$, and in Case II it is not. In the remaining three cases (Cases III, IV, and V), the point v selected by the spoiler is not near any point in G_0 that was selected earlier. In Case III, we have $v \in W$, while in Cases IV and V, we have $v \notin W$. In Case IV, the point v is not near any point in G_1 that was selected earlier, while in Case V, the point v is near some point in G_1 that was selected earlier.

Case I. *$\kappa \leq \kappa_0$ and ρ is completely contained in $A \cup B$.*

Since by the induction assumption λ_j is α_j, β_j-perfect, and since ρ is completely contained in $A \cup B$, it follows that λ_j is the identity on ρ. Let δ contain the set of points in a path of minimal length (in $U(G_0)$) connecting ρ and v. Thus, $|\delta| \leq \kappa_0$. We define λ_{j+1} by letting its domain H_{j+1} be $H_j \cup \delta$, and defining λ_{j+1} to be the identity on δ. We define the new partition P by replacing the class ρ by $\rho \cup \delta$. Since $|\delta| \leq \kappa_0 < \beta_j < \beta_0 < m$, it follows that $\rho \cup \delta \subseteq X$, where X is as above. The identity map on $\rho \cup \delta$ is then an isomorphism, because X does not contain the deleted forward edge e, and because G_0 and G_1 have the same d, q-coloring. Since λ_j is a partial d, q_j-isomorphism, it follows that λ_{j+1} is also a partial d, q_j-isomorphism, and hence a partial d, q_{j+1}-isomorphism (since $q_{j+1} < q_j$). Let α_{j+1} be $f^{\lfloor i/2 \rfloor + 1}(\alpha_j)$, and let β_{j+1} be $f^{i-1}(\alpha_j)$. We now show that condition 2(a) above holds, where j is replaced by $j + 1$. The first two inequalities are immediate. The third inequality holds, since

$$f^{3^{r-(j+1)+3}}(\alpha_{j+1}) = f^{i/3}(\alpha_{j+1}) = f^{(i/3)+\lfloor i/2 \rfloor+1}(\alpha_j),$$

which is less than $\beta_{j+1} = f^{i-1}(\alpha_j)$, since i is big enough that $(i/3) + \lfloor i/2 \rfloor + 1 < i - 1$. The final inequality holds, since

$$\beta_{j+1} = f^{i-1}(\alpha_j) = f^{3^{r-j+3}-1}(\alpha_j) < f^{3^{r-j+3}}(\alpha_j) < \beta_j.$$

As for condition 2(b), we see that

$$|H_{j+1}| = |H_j \cup \delta| \leq |H_j| + |\delta| \leq \alpha_j + \kappa_0 = \alpha_j + f^{\lfloor i/2 \rfloor}(\alpha_j) < f^{\lfloor i/2 \rfloor + 1}(\alpha_j) = \alpha_{j+1},$$

where the inequality follows from the growth rate of f (if we define z to be $f^{\lfloor i/2 \rfloor - 1}(\alpha_j)$, then $\alpha_j + f^{\lfloor i/2 \rfloor}(\alpha_j) < z + f(z) < 2f(z) = 2f^{\lfloor i/2 \rfloor}(\alpha_j) < f^{\lfloor i/2 \rfloor + 1}(\alpha_j)$). We now show that λ_{j+1} is α_{j+1}, β_{j+1}-perfect. We first show that P is an α_{j+1}, β_{j+1}-decomposition of H_{j+1}. To show this, we begin by showing that every pair of points within a class ρ' in P are within distance α_{j+1} of each other. If $\rho' \neq \rho \cup \delta$, then this follows from our induction assumption and the fact that $\alpha_j < \alpha_{j+1}$. If $\rho' = \rho \cup \delta$, then the distance between a pair of points in ρ' is at most $\alpha_j + \kappa_0$, which we showed is less than α_{j+1}. We now show that every pair of points that are in different classes in P are further than β_{j+1} apart. Since $\beta_{j+1} < \beta_j$, this follows from the inductive assumption unless one of the classes is $\rho \cup \delta$. But in this case, the distance is at least $\beta_j - \kappa_0 > f^i(\alpha_j) - f^{\lfloor i/2 \rfloor}(\alpha_j) > f^{i-1}(\alpha_j) = \beta_{j+1}$. So P is an α_{j+1}, β_{j+1}-decomposition of G_0. By a similar argument, $\lambda_{j+1}(P)$ is an α_{j+1}, β_{j+1}-decomposition of G_1. This establishes the first two conditions for λ_{j+1} to be α_{j+1}, β_{j+1}-perfect. The third condition holds for λ_{j+1} by construction, since it holds for λ_j.

Case II. $\kappa \leq \kappa_0$ *and* ρ *contains a point from* W.

Let δ contain those points in a path $v_0, \ldots, v_\kappa$ of minimal length (in $U(G_0)$) connecting ρ and v, where $v_0 \in \rho$ and $v_\kappa = v$. The domain H_{j+1} of λ_{j+1} will be the set $H_j \cup \delta$. We define $\lambda_{j+1}(v_\mu)$ by induction on μ. We describe only the case $\mu = 1$, since this is essentially the same as the general case.

Now the distance between A and W is greater than m, by definition of an m-trisection. The distance between an arbitrary pair of points in ρ is at most α_j, so the distance between an arbitrary pair of points in $\rho \cup \delta$ is at most $\alpha_j + \kappa_0$, which as we saw above is less than $f^{\lfloor i/2 \rfloor + 1}(\alpha_j)$, which in turn is less than $\beta_j < \beta_0 < m$. So, since $\rho \cap W \neq \varnothing$, it then follows that $(\rho \cup \delta) \cap A = \varnothing$. We know that there is an edge in $U(G_0)$ between v_0 and v_1. Assume for definiteness that, say, there is a directed edge $\langle v_1, v_0 \rangle$ but no directed edge $\langle v_0, v_1 \rangle$ in G_0. Let the $d, q_j - 1$-color of v_1 be Ω, and let the C_0-color of the edge $\langle v_1, v_0 \rangle$ be Ω'. Since v_0 and $\lambda_j(v_0)$ have the same d, q_j-color, and since $d = \beta_0 > \alpha_j \geq |H_j|$, we have that the number of points in G_0 with $d, q_j - 1$-color Ω and that have an edge in G_0 with C_0-color Ω' to v_0 but no reverse edge is the same as the number of points in G_1 with $d, q_j - 1$-color Ω and that have an edge in G_1 with C_0-color Ω' to $\lambda_j(v_0)$ and no reverse edge. Thus, there is a point w in G_1, distinct from any of $b_1, \ldots, b_j$, with $d, q_j - 1$-color Ω and that has an edge in G_1 with C_0-color Ω' to $\lambda_j(v_0)$, and no reverse edge. Let $\lambda_{j+1}(v_1) = w$. We define α_{j+1} and β_{j+1} in the same way as in Case I. By the same argument as before, conditions 2(a) and 2(b) above hold. The new partition P is obtained by replacing the class ρ by $\rho' = \rho \cup \delta$. As before, the first two conditions that λ_{j+1} be α_{j+1}, β_{j+1}-perfect hold. As for the third condition, we know that ρ, and hence ρ', is not contained in $A \cup B$, by assumption. If $\lambda_{j+1}(\rho')$, and hence $\lambda_j(\rho)$, were contained in $A \cup B$, then the restriction of λ_j onto ρ would be the identity, since λ_j is α_j, β_j-perfect. But then, since $\rho \cap W \neq \varnothing$, we would have $\lambda_j(\rho) \cap W \neq \varnothing$, a contradiction since $\lambda_j(\rho)$ is a subset of $A \cup B$. So λ_{j+1} is α_{j+1}, β_{j+1}-perfect, as desired.

We now show that λ_{j+1} is a partial d, q_{j+1}-isomorphism. By the induction assumption, we need only show that the restriction of λ_{j+1} to ρ' is a d, q_{j+1}-isomorphism. Since λ_j is a partial d, q_j-isomorphism by the induction assumption, and since $q_j - q_{j+1} > \kappa_0 \geq \kappa$, it is not hard to see that our construction guarantees that the restriction of λ_{j+1} to ρ' is a d, q_{j+1}-isomorphism, provided that both ρ' and $\lambda_{j+1}(\rho')$ contain no cycles. This is because we extend the isomorphism by "following the path", and, intuitively, we never get into trouble because there are no cycles. The straightforward details are left to the reader. So we now show that both ρ' and $\lambda_{j+1}(\rho')$ contain no cycles. As we showed, $\rho' \cap A = \emptyset$, where we simply used the fact that ρ' contains a point in W, and ρ' is a connected set whose size is less than m. By a similar argument, $\lambda_{j+1}(\rho') \cap A = \emptyset$. Since $|\rho'| < m$, it follows that if ρ' were to contain a cycle, it would be of size less than m. By definition of an m-trisection, every point on a cycle shorter than m is contained in A. Since $\rho' \cap A = \emptyset$, it follows that ρ' contains no cycles. Similarly, neither does $\lambda_{j+1}(\rho')$.

Case III. $\kappa > \kappa_0$ *and* $v \in W$.

If the G_1-distance between v and $\lambda_j(H_j)$ is greater than κ_0, then let $\lambda_{j+1}(v) = v$. So assume for now that the G_1-distance (and hence the G_0-distance) between v and $\lambda_j(H_j)$ is at most κ_0.

Let Ω be the d, q_j-color of v. By the definition of an m-trisection, we know that Ω is m-dispersed in W. So there is a subset R of W, consisting entirely of points of color Ω, so that $|R| = m$ and so that every distinct pair of points in R is of G_0-distance greater than m apart. Now each point in $\lambda_j(H_j)$ is within G_0-distance κ_0 of at most one member of R, since otherwise two distinct points in R would be within G_0-distance $2\kappa_0 + 1 < m$ of each other. Since $|\lambda_j(H_j)| \leq \alpha_j < m$, it follows that there is some member w of R that is of G_0-distance (and hence G_1-distance) greater than κ_0 from every member of $\lambda_j(H_j)$. Let $\lambda_{j+1}(v) = w$.

So in each situation (whether or not the G_1-distance between v and $\lambda_j(H_j)$ is greater than κ_0), we have defined $\lambda_{j+1}(v)$ so that the G_1-distance between $\lambda_{j+1}(v)$ and $\lambda_j(H_j)$ is greater than κ_0. Let $\alpha_{j+1} = \alpha_j + 1$, and let $\beta_{j+1} = \kappa_0$. It is easy to verify that the induction assumption holds when j is replaced by $j + 1$.

Case IV. $\kappa > \kappa_0$ *and* $v \in A \cup B$, *and the* G_1*-distance between* v *and* $\lambda_j(H)$ *is greater than* κ_0.

Let $\lambda_{j+1}(v) = v$, and let $\alpha_{j+1} = \alpha_j + 1$ and $\beta_{j+1} = \kappa_0$. Again, it is easy to verify that the induction assumption holds when j is replaced by $j + 1$.

Case V. $\kappa > \kappa_0$ *and* $v \in A \cup B$, *and the* G_1*-distance between* v *and* $\lambda_j(H_j)$ *is at most* κ_0.

We shall first show that there is a point $w \in W$ where $d_{G_0}(w, v) < m$. Let $\lambda_j(\rho')$ be the closest class (in $U(G_1)$) to v in $\lambda_j(H_j)$, and let z be a point in $\lambda_j(\rho')$ whose G_1-distance from v is at most κ_0. Now ρ' and $\lambda_j(\rho')$ are distinct, because the G_0-distance (and hence the G_1-distance) between v and ρ' is more than κ_0, while the G_1-distance between v and $\lambda_j(\rho')$ is at most κ_0. So, from the third condition on λ_j being α_j, β_j-perfect, we know that $\lambda_j(\rho')$ has a point w in W. Now $\lambda_j(\rho')$ is connected, since ρ' is and since λ_j is a partial isomorphism. Therefore, there is a path between z and w completely within $\lambda_j(\rho')$, since both are in $\lambda_j(\rho')$. The shortest such path is of length at most α_j, since $|\lambda_j(\rho')| = |\rho'| \leq |H_j| \leq \alpha_j$. Since the G_1-distance, and hence the G_0-distance between v and z is at most κ_0, and since, as we just saw, the G_0-distance between z and w is at most α_j, it follows that the G_0-distance between v and w is at most $\kappa_0 + \alpha_j$. As we showed before, $\kappa_0 + \alpha_j < m$. So there is a point $w \in W$ where $d_{G_0}(w, v) < m$, as claimed.

Since $w \in W$, we can now apply Case I, II, or III to extend λ_j so that the domain includes w, just as if the spoiler had selected w in G_0 as his move (Cases I, II, and III are the only cases where the spoiler selects a point in W). Since the distance between w and v is less than m, and since the G_0-distance between A and W is greater than m, it follows that the minimal path in $U(G_0)$ from w to v contains no point in A. Therefore, by the same argument as we gave in Case II, we can further extend λ_j to the minimal path in $U(G_0)$ connecting w and v, to obtain λ_{j+1}.

This completes the induction step, and hence the proof.

Acknowledgments. The authors are grateful to Phokion Kolaitis, Moshe Vardi, Larry Stockmeyer, Joe Halpern, Cynthia Dwork, Janos Simon, and Neil Immerman for helpful comments. Special thanks go to Stavros Cosmadakis and Yuri Gurevich for a careful reading of the paper.

REFERENCES

[AA] A. Aggarwal and R. J. Anderson, *A random NC algorithm for depth-first search*, ***Proceedings of the Nineteenth ACM Symposium on Theory of Computing*** (1987), pp. 325–334.

[Aj] M. Ajtai, Σ_1^1*-formulae on finite structures*, ***Annals of Pure and Applied Logic***, vol. 24 (1983), pp. 1–48.

[AKLLR] R. Aleliunas, R. M. Karp, R. J. Lipton, L. Lovász, and C. Rackoff, *Random walks, universal traversal sequences, and the complexity of maze problems*, ***Proceedings of the Twentieth IEEE Symposium on Foundations of Computer Science*** (1979), pp. 218–223.

[An] R. Anderson, Private communication, May 1987.

[BKBR] C. Beeri, P. Kanellakis, F. Bancilhon, and R. Ramakrishnan, *Bounds on the propagation of selection into logic programs*, ***Proceedings of the Sixth ACM Symposium on Principles of Database Systems*** (1987), pp. 214–226. (Final version, to appear in ***Journal of Computer and System Sciences***.)

[BeSi] P. Berman and J. Simon, *Lower bounds on graph threading by probabilistic machines*, ***Proceedings of the Twenty-fourth IEEE Symposium on Foundations of Computer Science*** (1983), pp. 304–311.

[Bo] B. Bollobás, ***Random graphs***, Academic Press, New York, 1982.

[BoSi] R. B. Boppana and M. Sipser, *The complexity of finite functions*, ***Handbook of Theoretical Computer Science*** (to appear).

[BCDRT] A. Borodin, S. A. Cook, P. W. Dymond, W. L. Ruzzo, and M. Tompa, *Two applications of complementation via inductive counting*, ***Proceedings of the Third Conference on Structure in Complexity Theory*** (1988), pp. 116–125. (Final version, *Two applications of inductive counting for complementation problems*, ***SIAM Journal on Computing***, vol. 18 (1989), pp. 559–578.)

[CH] A. K. Chandra and D. Harel, *Structure and complexity of relational queries*, ***Journal of Computer and System Sciences***, vol. 25 (1982), 99–128.

[Che] H. Chernoff, *A measure of asymptotic efficiency for tests of hypothesis based on the sum of observations*, ***Annals of Mathematical Statistics***, vol. 23 (1952), pp. 493–507.

[Chv] V. Chvátal, ***On the computational complexity of finding a kernel***, Report No. CRM-300, Centre de Recherches Mathématiques, Université de Montréal, Montréal, 1973.

[CV] R. Cole and U. Vishkin, *Approximate and exact parallel scheduling with applications to list, tree and graph problems*, ***Proceedings of the Twenty-seventh IEEE Symposium on Foundations of Computer Science*** (1986), pp. 478–491.

[CR] S. A. Cook and C. W. Rackoff, *Space lower bounds for maze threadability on restricted machines*, ***SIAM Journal on Computing***, vol. 9 (1980), pp. 636–652.

[CW] D. Coppersmith and S. Winograd, *Matrix multiplication via arithmetic progressions*, ***Proceedings of the Nineteenth ACM Symposium on Theory of Computing*** (1987), pp. 1–6.

[Eh] A. Ehrenfeucht, *An application of games to the completeness problem for formalized theories*, ***Fundamenta Mathematicae***, vol. 49 (1961), pp. 129–141.

[Fa1] R. Fagin, *Generalized first-order spectra and polynomial-time recognizable sets*, ***Complexity of computation*** (R. Karp, editor), SIAM-AMS Proceedings, vol. 7, 1974, pp. 43–73.

[Fa2] ———, *Monadic generalized spectra*, ***Zeitschrift für Mathematische Logik und Grundlagen der Mathematik***, vol. 21 (1975), pp. 89–96.

[Fr] R. Fraïssé, *Sur quelques classifications des systèmes de relations*, ***Publications Scientifiques de l'Université d'Alger***, vol. 1 (1954), pp. 35–182.

[GJ] M. GAREY and D. S. JOHNSON, ***Computers and intractibility: a guide to the theory of NP-completeness***, Freeman, San Francisco, California, 1979.

[HMS] P. H. HOCHSCHILD, E. W. MAYR, and A. R. SIEGEL, *Parallel Graph Algorithms*, Stanford University Technical Report, 1987; to appear in ***SIAM Journal on Computing***.

[Im1] N. IMMERMAN, *Upper and lower bounds for first-order expressibility*, ***Journal of Computer and System Sciences***, vol. 25 (1982), pp. 76–98.

[Im2] ———, *Relational queries computable in polynomial time*, ***Information and Control***, vol. 68 (1986), pp. 86–104.

[Im3] ———, *Languages which capture complexity classes*, ***SIAM Journal on Computing***, vol. 16 (1987), pp. 760–778.

[Im4] ———, *Descriptive and computational complexity*, ***Computational complexity theory***, Proceedings of Symposia in Applied Mathematics, vol. 38, American Mathematical Society, Providence, Rhode Island, 1989, pp. 75–91.

[Ka1] P. KANELLAKIS, Private communication, October 1986.

[Ka2] ———, *Logic programming and parallel complexity*, ***Foundations of deductive databases and logic programming*** (J. Minker, editor), Morgan Kaufmann, Los Altos, California, 1988.

[KW] M. KARCHMER and A. WIGDERSON, *Monotone circuits for connectivity require superlogarithmic depth*, ***Proceedings of the Twentieth ACM Symposium on Theory of Computing*** (1988), pp. 539–550.

[KV] P. G. KOLAITIS and M. Y. VARDI, *The decision problem for the probabilities of higher-order properties*, ***Proceedings of the Nineteenth ACM Symposium on Theory of Computing*** (1987), pp. 425–435.

[Le] D. LEIVANT, *Characterization of complexity classes in higher-order logic*, ***Proceedings of the Second IEEE Conferencee on Structure in Complexity Theory*** (1987), pp. 203–217. (Revised version, *Descriptive characterizations of computational complexity*, ***Journal of Computer and System Sciences***, vol. 39 (1989), pp. 51–83.)

[Lo] L. LOVÁSZ, *Computing ears and branchings in parallel*, ***Proceedings of the Twenty-sixth IEEE Symposium on Foundations of Computer Science*** (1985), pp. 464–467.

[Ly] J. LYNCH, *Complexity classes and theories of finite models*, ***Mathematical Systems Theory***, vol. 15 (1982), pp. 127–144.

[PR] V. PAN and J. REIF, *Efficient parallel solution of linear systems*, ***Proceedings of the Seventeenth ACM Symposium on Theory of Computing*** (1985), pp. 143–152.

[Ro] M. DE ROUGEMONT, *Second-order and inductive definability on finite structures*, ***Zeitschrift für Mathematische Logik und Grundlagen der Mathematik***, vol. 33 (1987), pp. 47–63. (Preliminary version appeared as: *Uniform definability on structures with successor*, ***Proceedings of the Sixteenth ACM Symposium on Theory of Computing*** (1984), pp. 409–417.)

[Sa] W. J. SAVITCH, *Relations between nondeterministic and deterministic tape complexities*, ***Journal of Computer and System Sciences***, vol. 4 (1970), pp. 177–192.

[Ul] J. D. ULLMAN, *Database theory: past and future*, ***Proceedings of the Sixth ACM Symposium on Principles of Database Systems*** (1987), pp. 1–10.

[Va] M. Y. VARDI, *The complexity of relational query languages*, ***Proceedings of the Fourteenth ACM Symposium on Theory of Computing*** (1982), pp. 137–146.

IBM ALMADEN RESEARCH CENTER
SAN JOSE, CALIFORNIA 95120