

UNARY AUTOMATIC GRAPHS: ALGORITHMIC PERSPECTIVE

BAKHADYR KHOUSSAINOV, JIAMOU LIU, AND MIA MINNES

ABSTRACT. This paper studies infinite graphs produced from a natural unfolding operation applied to finite graphs. Graphs produced via such operations are of finite degree and automatic over the unary alphabet (that is, they can be described by finite automata over unary alphabet). We investigate algorithmic properties of such unfolded graphs given their finite presentations. In particular, we ask whether a given node belongs to an infinite component, whether two given nodes in the graph are reachable from one another, and whether the graph is connected. We give polynomial time algorithms for each of these questions. For a fixed input graph, the algorithm for the first question is in constant time and the second question is decided using an automaton that recognizes reachability relation in a uniform way. Hence, we improve on previous work, in which non-elementary or non-uniform algorithms were found.

1. INTRODUCTION

We study the algorithmic properties of infinite graphs that result from a natural unfolding operation applied to finite graphs. The unfolding process always produces infinite graphs of finite degree. Moreover, the class of resulting graphs is a subclass of the class of automatic graphs. As such, any element of this class possesses all the known algorithmic and algebraic properties of automatic structures. An equivalent way to describe these graphs employs automata over a unary alphabet (see Theorem 4.4). Therefore, we call this class of graphs *unary automatic graphs of finite degree*.

In recent years there has been increasing interest in the study of structures that can be presented by automata. The underlying idea in this line of research consists of using automata (such as word automata, Büchi automata, tree automata, or Rabin automata) to represent structures and study logical and algorithmic consequences of such presentations. Informally, a structure $\mathcal{A} = (A; R_0, \dots, R_m)$ is *automatic* if the domain A and all the relations R_0, \dots, R_m of the structure are recognized by finite automata (precise definitions are in the next section). For instance, an automatic graph is one whose set of vertices and set of edges can each be recognized by finite automata. Automatic structures possess a number of nice algorithmic and model-theoretic properties. For example, Khoussainov and Nerode proved that the first-order theory of any automatic structure is decidable [9]. This result is extended by adding the \exists^∞ (there are infinitely many) and $\exists^{n,m}$ (there are m many mod n) quantifiers to the first order logic (see [2], [12]). Blumensath and Grädel proved a logical characterization theorem stating that automatic structures are exactly those definable in the following fragment of the arithmetic $(\omega; +, |_2, \leq, 0)$, where $+$ and \leq have their usual meanings and $|_2$ is a weak divisibility predicate for which $x|_2y$ if and only if x is a power of 2 and divides y (see [2]). Automatic structures are closed under first-order interpretations. There are descriptions of automatic linear orders and trees in terms of model theoretic concepts such as Cantor-Bendixson ranks [17]. Also, Khoussainov, Nies, Rubin and Stephan have characterized the isomorphism types of automatic Boolean algebras [10]; Thomas and Oliver have given a full description of finitely generated automatic groups [15]. Some of these results have direct algorithmic implications. For example, isomorphism problem for automatic well-ordered sets and Boolean algebras is decidable [17].

Date: December 17, 2007.

Key words and phrases. automatic structures, automata, connected components of graphs, graph algorithms.

There is also a body of work devoted to the study of resource-bounded complexity of the first order theories of automatic structures. For example, on the one hand, Grädel and Blumensath constructed examples of automatic structures whose first-order theories are non-elementary (see [2]). On the other hand, Kuske and Lohrey in [14] proved that the first-order theory of any automatic graph of bounded degree is elementary. It is worth noting that when both a first-order formula and an automatic structure \mathcal{A} are fixed, determining if a tuple \bar{a} from \mathcal{A} satisfies $\phi(\bar{x})$ can be done in linear time.

Most of the results about automatic structures, including the ones mentioned above, demonstrate that in various concrete senses automatic structures are not complex from a logical point of view. However, this intuition can be misleading. For example, in [10] it is shown that the isomorphism problem for automatic structures is Σ_1^1 -complete. This informally tells us that there is no hope for a description (in a natural logical language) of the isomorphism types of automatic structures. Also, Khoussainov and Minnes [8] provide examples of automatic structures whose Scott ranks can be as high as possible, fully covering the interval $[1, \omega_1^{CK} + 1]$ of ordinals (where ω_1^{CK} is the first non-computable ordinal). They also show that the ordinal heights of well-founded automatic relations can be arbitrarily large ordinals below ω_1^{CK} .

In this paper, we study the class of unary automatic graphs of finite degree. Since these graphs are described by the unfolding operation (Definition 4.3) on the pair of finite graphs $(\mathcal{D}, \mathcal{F})$, we use this pair to represent the graph. The size of this pair is the sum of the sizes of the automata that represent these graphs. In the study of algorithmic properties of these graphs one directly deals with the pair $(\mathcal{D}, \mathcal{F})$. We are interested in the following natural decision problems:

- **Connectivity Problem.** Given an automatic graph \mathcal{G} , decide if \mathcal{G} is connected.
- **Reachability Problem.** Given an automatic graph \mathcal{G} and two vertices x and y of the graph, decide if there is a path from x to y .

If we restrict to the class of finite graphs, these two problems are decidable and can be solved in linear time on the sizes of the graphs. However, we are interested in infinite graphs and therefore much more work is needed to investigate the problems above. In addition, we also pose the following two problems:

- **Infinity Testing Problem.** Given an automatic graph \mathcal{G} and a vertex x , decide if the component of \mathcal{G} containing x is infinite.
- **Infinite Component Problem.** Given an automatic graph \mathcal{G} decide if \mathcal{G} has an infinite component.

Unfortunately, for the class of automatic graphs all of the above problems are undecidable. In fact, one can provide exact bounds on this undecidability. The connectivity problem is Π_2^0 -complete; the reachability problem is Σ_1^0 -complete; the infinite component problem is Σ_3^0 -complete; and the infinity testing problem is Π_2^0 -complete (see [17]).

Since all unary automatic structures are first-order definable in $S1S$ (the monadic second-order logic of the successor function), it is not hard to prove that all the problems above are decidable ([1], [17]). Direct constructions using this definability in $S1S$ yield algorithms with non-elementary time since one needs to transform $S1S$ formulas into automata ([4]). However, we provide polynomial time algorithms for solving all the above problems for this class of graphs. We now outline the rest of this paper by explaining the main results.

Section 2 introduces the main definitions needed, including the concept of automatic structure. Section 3 singles out unary automatic graphs and provides a characterization theorem (Theorem 3.4). That section also gives a direct proof that unary automatic graphs have prefix-free presentations. Section 4 introduces unary automatic graphs of finite degree. The main result is Theorem 4.4 that explicitly provides an algorithm for building unary automatic graphs of finite degree. This theorem is used throughout the paper. Section 5 is devoted to deciding the infinite component problem. The main result is the following:

Theorem 5.1 *The infinite component problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(n^{\frac{3}{2}})$, where n is the number of states of the unary automaton recognizing \mathcal{G} .*

In this section, we make use of the concept of oriented walk for finite directed graphs. The subsequent section is devoted to deciding the infinity testing problem. The main result is the following:

Theorem 6.1 *The infinity testing problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(n^{\frac{5}{2}})$, where n is the number of states of the unary automaton \mathcal{A} recognizing \mathcal{G} . In particular, when \mathcal{A} is fixed, there is a constant time algorithm that decides the infinity testing problem on \mathcal{G} .*

The fact that there is a constant time algorithm when \mathcal{A} is fixed will be made clear in the proof. The value of the constant is polynomial in the number of states of \mathcal{A} .

The reachability problem is addressed in Section 7. This problem has been studied in [3], [6], [18] via the class of **pushdown graphs**. A pushdown graph is the configuration space of a pushdown automaton. Unary automatic graphs are pushdown graphs [18]. In [3], [6], [18] it is proved that for a pushdown graph G , given a node v , there is an automaton that recognizes all nodes reachable from v . The number of states in the automaton depends on the input node v . This result implies that there is an algorithm that decides the reachability problem on unary automatic graphs of finite degree. However, there are several issues with this algorithm. The automata constructed by the algorithm are not uniform in v in the sense that different automata are built for different input nodes v . Moreover, the automata are nondeterministic. Hence, the size of the deterministic equivalent automata is exponential in the size of the representation of v . Section 7 provides an alternative algorithm to solve the reachability problem on unary automatic graphs of finite degree. This new algorithm constructs a deterministic automaton \mathcal{A}_{Reach} that accepts the set of pairs $\{(u, v) \mid \text{there is a path from } u \text{ to } v\}$. The size of \mathcal{A}_{Reach} only depends on the number of states of the automaton n , and constructing the automaton requires polynomial time in n . The practical advantage of such a uniform solution is that, when \mathcal{A}_{Reach} is built, deciding whether node v is reachable from u by a path takes only linear time (details are in Section 7). The main result of this section is the following:

Theorem 7.1 *Suppose \mathcal{G} is a unary automatic graph of finite degree represented by unary automaton \mathcal{A} of size n . There exists a polynomial-time algorithm that solves the reachability problem on \mathcal{G} . For inputs u, v , the running time of the algorithm is $O(|u| + |v| + n^{\frac{5}{2}})$.*

Finally, the last result of this section follows from the algorithm for the reachability problem and solves the connectivity problem for G .

Theorem 7.8 *The connectivity problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(n^3)$, where n is the number of states of the unary automaton recognizing \mathcal{G} .*

2. PRELIMINARIES

A **finite automaton** \mathcal{A} over an alphabet Σ is a tuple (S, ι, Δ, F) , where S is a finite set of **states**, $\iota \in S$ is the **initial state**, $\Delta \subset S \times \Sigma \times S$ is the **transition table** and $F \subset S$ is the set of **final states**. A **computation** of \mathcal{A} on a word $\sigma_1\sigma_2 \dots \sigma_n$ ($\sigma_i \in \Sigma$) is a sequence of states, say q_0, q_1, \dots, q_n , such that $q_0 = \iota$ and $(q_i, \sigma_{i+1}, q_{i+1}) \in \Delta$ for all $i \in \{0, 1, \dots, n-1\}$. If $q_n \in F$, then the computation is **successful** and we say that automaton \mathcal{A} *accepts* the word. The **language** accepted by the automaton \mathcal{A} is the set of all words accepted by \mathcal{A} . In general, $D \subset \Sigma^*$ is **FA recognizable**, or **regular**, if D is the language accepted by some finite automaton. In this paper we always assume the automata are deterministic. For two states q_0, q_1 , the **distance** from q_0 to q_1 is the minimum number of transitions required for \mathcal{A} to go from q_0 to q_1 .

To formalize the notion of a relation being recognized by an automaton, we define synchronous n -tape automata. Such an automaton can be thought of as a one-way Turing machine with n input tapes. Each tape is semi-infinite having written on it a word in the alphabet Σ followed by a succession of \diamond symbols. The automaton starts in the initial state, reads simultaneously the first

symbol of each tape, changes state, reads simultaneously the second symbol of each tape, changes state, etc., until it reads \diamond on each tape. The automaton then stops and accepts the n -tuple of words if and only if it is in a final state.

More formally, we write Σ_\diamond for $\Sigma \cup \{\diamond\}$ where \diamond is a symbol not in Σ . The **convolution** of a tuple $(w_1, \dots, w_n) \in \Sigma^{*n}$ is the string $\otimes(w_1, \dots, w_n)$ of length $\max_i |w_i|$ over the alphabet $(\Sigma_\diamond)^n$ which is defined as follows: the k^{th} symbol is $(\sigma_1, \dots, \sigma_n)$ where σ_i is the k^{th} symbol of w_i if $k \leq |w_i|$, and is \diamond otherwise. The **convolution** of a relation $R \subset \Sigma^{*n}$ is the relation $\otimes R \subset (\Sigma_\diamond)^{n*}$ formed as the set of convolutions of all the tuples in R . An n -**tape automaton** on Σ is a finite automaton over the alphabet $(\Sigma_\diamond)^n$. An n -ary relation $R \subset \Sigma^{*n}$ is **FA recognizable**, or **regular**, if its convolution $\otimes R$ is recognizable by an n -tape automaton.

A **structure** \mathcal{S} consists of a countable set D called the **domain** and some relations and operations on D . We may assume that \mathcal{S} only contains relational predicates since any operations can be replaced with their graphs. We write $\mathcal{S} = (D, R_1^D, \dots, R_k^D, \dots)$ where R_i^D is an n_i -ary relation on D . The relation R_i are sometimes called basic or atomic relations. We assume that the function $i \mapsto n_i$ is always a computable one. For this paper we consider infinite structures with finitely many relations.

A structure \mathcal{S} is **automatic over alphabet** Σ if its domain $D \subset \Sigma^*$ is finite automaton recognizable, and there is an algorithm that for each i produces an n_i -tape automaton recognizing the relation $R_i^D \subset (\Sigma^*)^{n_i}$. A structure is called **automatic** if it is automatic over some alphabet. If \mathcal{B} is isomorphic to an automatic structure \mathcal{S} , then we call \mathcal{S} an **automatic presentation** of \mathcal{B} and say that \mathcal{B} is **automatically presentable**.

An example of an automatic structure is the word structure $(\{0, 1\}^*, L, R, E, \preceq)$, where for all $x, y \in \{0, 1\}^*$, $L(x) = x0$, $R(x) = x1$, $E(x, y)$ if and only if $|x| = |y|$, and \preceq is the lexicographical order. The configuration graph of any Turing machine is another example of an automatic structure. Examples of automatically presentable structures are $(\mathbb{N}, +)$, (\mathbb{N}, \leq) , (\mathbb{N}, S) , the group $(\mathbb{Z}, +)$, the order on the rational (\mathbb{Q}, \leq) , and the Boolean algebra of finite and co-finite subsets of \mathbb{N} . Consider the first-order logic extended by \exists^ω (there exist infinitely many) and $\exists^{n,m}$ (there exist n many mod m , where n and m are natural numbers) quantifiers. We denote this logic by $FO + \exists^\omega + \exists^{n,m}$. We will use the following theorem without explicit reference to it.

Theorem 2.1. [9] *Let \mathcal{A} be an automatic structure. There exists an algorithm that, given a formula $\phi(\bar{x})$ in $FO + \exists^\omega + \exists^{n,m}$, produces an automaton that recognizes exactly those tuples \bar{a} from the structure that make ϕ true. In particular, the set of all sentences of $FO + \exists^\omega + \exists^{n,m}$ which are true in \mathcal{A} is decidable.*

3. UNARY AUTOMATIC GRAPHS

We now turn our attention to the subclass of the automatic structures which is the focus of the paper.

Definition 3.1. *A structure \mathcal{A} is **unary automatic** if it has an automatic presentation where all automata (those representing the domain and relations of \mathcal{A}) are over the unary alphabet $\{1\}$.*

Examples of unary automatic structures are (ω, S) and (ω, \leq) . Some recent work on unary automatic structures includes a characterization of unary automatic linearly ordered sets, permutation structures, graphs, and equivalence structures (see [11], [1]). For example, unary automatic linearly ordered sets are exactly those that are isomorphic to a finite sum of orders of type ω , ω^* (the order of negative integers), and finite n .

Definition 3.2. *A **unary automatic graph** is a graph (V, E) whose domain is $\{1\}^*$, and whose edge relation E is regular.*

Convention. To eliminate bulky exposition, we make the following assumptions in the rest of the paper.

- We restrict our attention to unary automatic structures. Hence, when we write “automatic structure”, we mean “unary automatic structure”.
- All structures are infinite unless explicitly specified.
- The domains of unary automatic structures coincide with the set 1^* of all unary strings $\{\lambda, 1, 11, 111, \dots\}$. All the notions and results below can be adapted to the case when the domains are regular subsets of 1^* ¹.
- The graphs are undirected. The case of directed graphs can be treated in a similar manner.




FIGURE 1. A Typical Unary Graph Automaton

Let $G = (V, E)$ be an automatic graph. Let \mathcal{A} be an automaton recognizing E . We establish some terminology for the automaton \mathcal{A} . The general shape of \mathcal{A} is given in Figure 1. All the states reachable from the initial state by reading inputs of type $(1, 1)$ are called $(1, 1)$ -states. A **tail** in \mathcal{A} is a sequence of states linked by transitions without repetition. A **loop** is a sequence of states linked by transitions such that the last state coincides with the first one, and with no repetition in the middle. The set of $(1, 1)$ -states is a disjoint union of a tail and a loop. We call the tail the $(1, 1)$ -tail and the loop the $(1, 1)$ -loop. Let s be a $(1, 1)$ state. All the states reachable from s by reading inputs of type $(1, \diamond)$ are called $(1, \diamond)$ -states. This collection of all $(1, \diamond)$ -states is also a disjoint union of a tail and a loop (see the figure), called the $(1, \diamond)$ -tail and the $(1, \diamond)$ -loop, respectively. The $(\diamond, 1)$ -tails and $(\diamond, 1)$ -loops are defined in a similar matter.

We say that an automaton is **standard** if the lengths of all its loops and tails equal some number p , called the **loop constant**. If \mathcal{A} is standard unary automaton recognizing a binary relation, it has exactly $2p$ $(1, 1)$ -states. On each of these states, there is a $(1, \diamond)$ -tail and a $(\diamond, 1)$ -tail of length exactly p . At the end of each $(1, \diamond)$ -tail and $(\diamond, 1)$ -tail there is a $(1, \diamond)$ -loop and $(\diamond, 1)$ -loop, respectively, of size exactly p . Therefore If n is the number of states in \mathcal{A} , then $n = 8p^2$.

Lemma 3.3. *Let \mathcal{A} be an n state automaton recognizing a binary relation E on 1^* . There exists an equivalent standard unary automaton with at most $8n^{2n}$ states.*

Proof. Let p be the least common multiple of the lengths of all loops and tails of \mathcal{A} . An easy estimate shows that p is no more than n^n . One can transform \mathcal{A} into an equivalent standard automaton whose loop constant is p . Hence, there is a standard unary automaton equivalent to \mathcal{A} whose size is bounded above by $8n^{2n}$. \square

We can simplify the general shape of the automaton using the fact that we consider undirected graphs. Indeed, we need only consider transitions labelled by $(\diamond, 1)$. To see this, given an automaton

¹A language $L \subseteq 1^*$ is regular if and only if L is a union of a finite languages together with a language of a fixed period. This observation should be used for the cases when the domains of the structures do not coincide with 1^*

with only $(\diamond, 1)$ transitions, to include all symmetric transitions, add a copy of each $(\diamond, 1)$ transition which is labelled with $(1, \diamond)$.

We recall a characterization theorem of unary automatic graphs from [17]. Let $\mathcal{B} = (B, E_B)$ and $\mathcal{D} = (D, E_D)$ be finite graphs. Let R_1, R_2 be subsets of $D \times B$, and R_3, R_4 be subsets of $B \times B$. Consider the graph \mathcal{D} followed by ω many copies of \mathcal{B} , ordered as $\mathcal{B}^0, \mathcal{B}^1, \mathcal{B}^2, \dots$. Formally, the vertex set of \mathcal{B}^i is $B \times \{i\}$ and we write $b^i = (b, i)$ for $b \in B$ and $i \in \omega$. The edge set E^i of \mathcal{B}^i consists of all pairs (a^i, b^i) such that $(a, b) \in E_B$. We define the infinite graph, $unwind(\mathcal{B}, \mathcal{D}, \bar{R})$, as follows: 1) The vertex set is $D \cup B^0 \cup B^1 \cup B^2 \cup \dots$; 2) The edge set contains $E_D \cup E^0 \cup E^1 \cup \dots$ as well as the following edges, for all $a, b \in B$, $d \in D$, and $i, j \in \omega$:

- (d, b^0) when $(d, b) \in R_1$, and (d, b^{i+1}) when $(d, b) \in R_2$,
- (a^i, b^{i+1}) when $(a, b) \in R_3$, and (a^i, b^{i+2+j}) when $(a, b) \in R_4$.

Theorem 3.4. [17] *A graph has a unary automaton presentation if and only if it is isomorphic to $unwind(\mathcal{B}, \mathcal{D}, \bar{R})$ for some parameters \mathcal{B} , \mathcal{D} , and \bar{R} . Moreover, if \mathcal{A} is a standard automaton representing \mathcal{G} then the parameters $\mathcal{B}, \mathcal{D}, \bar{R}$ can be extracted in $O(n^2)$; otherwise, the parameters can be extracted in $O(n^{2^n})$, where n is the number of states in \mathcal{A} .*

4. UNARY AUTOMATIC GRAPHS OF FINITE DEGREE

A graph is of **finite degree** if there are at most finitely many edges from each vertex v . We call an automaton \mathcal{A} recognizing a binary relation over $\{1\}$ a **one-loop automaton** if its transition diagram contains exactly one loop, the $(1, 1)$ -loop. The general structure of one-loop automata is given in Figure 2.




FIGURE 2. One-loop automaton

We will always assume that the lengths of all the tails of the one-loop automata are not bigger than the size of the $(1, 1)$ -loop. The following is an easy proposition and we omit its proof.

Proposition 4.1. *Let $\mathcal{G} = (V, E)$ be a unary automatic graph, then \mathcal{G} is of finite degree if and only if there is a one-loop unary automaton \mathcal{A} recognizing E . \square*

By Lemma 3.3, transforming a given unary automaton to an equivalent standard automaton may blow up the number of states exponentially. However, if \mathcal{A} is a one-loop automaton with n states, the equivalent standard one-loop automaton has its $(1, 1)$ -loop of size at most n , and hence has at most $4n^2$ states. Therefore we always assume the input automaton \mathcal{A} is standard. Let p be the loop constant of \mathcal{A} , then \mathcal{A} has exactly $4p^2$ states. In the rest of the paper, we will state all results in terms of p instead of n , the number of states of the input automaton.

Given two unary automatic graphs of finite degree $\mathcal{G}_1 = (V, E_1)$ and $\mathcal{G}_2 = (V, E_2)$ (where we recall the convention that the domain of each graph is 1^*), we can form the **union graph** $\mathcal{G}_1 \oplus \mathcal{G}_2 = (V, E_1 \cup E_2)$ and the **intersection graph** $\mathcal{G}_1 \otimes \mathcal{G}_2 = (V, E_1 \cap E_2)$. Automatic graphs of finite degree are closed under these operations. Indeed, let \mathcal{A}_1 and \mathcal{A}_2 be one-loop automata recognizing E_1 and E_2 with loop constants p_1 and p_2 , respectively. The standard construction that builds automata for the union and intersection operations produces a one-loop automaton whose loop constant is

$p_1 \cdot p_2$. We introduce another operation: consider the new graph $\mathcal{G}'_1 = (V, E'_1)$, where the set E'_1 of edges is defined as follows; a pair $(1^n, 1^m)$ is in E' if and only if $(1^n, 1^m) \notin E$ and $|n - m| \leq p_1$. The relation E'_1 is recognized by the same automaton as E_1 , modified so that all $(\diamond, 1)$ -states that are final declared non-final, and all the $(\diamond, 1)$ -states that are non-final declared final. Thus, we have the following proposition:

Proposition 4.2. *If \mathcal{G}_1 and \mathcal{G}_2 are automatic graphs of finite degree then so are $\mathcal{G}_1 \oplus \mathcal{G}_2$, $\mathcal{G}_1 \otimes \mathcal{G}_2$, and \mathcal{G}'_1 . \square*

Now our goal is to recast Theorem 3.4 for graphs of finite degree. Our analysis will show that, in contrast to the general case for automatic graphs, the parameters \mathcal{B} , \mathcal{D} , and \bar{R} for graphs of finite degree can be extracted in linear time.

Definition 4.3 (Unfolding Operation). *Let $\mathcal{D} = (V_{\mathcal{D}}, E_{\mathcal{D}})$ and $\mathcal{F} = (V_{\mathcal{F}}, E_{\mathcal{F}})$ be finite graphs. Consider the finite sets $\Sigma_{\mathcal{D}, \mathcal{F}}$ consisting all mappings $\eta : V_{\mathcal{D}} \rightarrow P(V_{\mathcal{F}})$, and $\Sigma_{\mathcal{F}}$ consisting of all mappings $\sigma : V_{\mathcal{F}} \rightarrow P(V_{\mathcal{F}})$. Any infinite sequence $\alpha = \eta\sigma_0\sigma_1 \dots$ where $\eta \in \Sigma_{\mathcal{D}, \mathcal{F}}$ and $\sigma_i \in \Sigma_{\mathcal{F}}$ for each i , defines the infinite graph $\mathcal{G}_{\alpha} = (V_{\alpha}, E_{\alpha})$ as follows:*

- $V_{\alpha} = V_{\mathcal{D}} \cup \{(v, i) \mid v \in V_{\mathcal{F}}, i \in \omega\}$.
- $E_{\alpha} = E_{\mathcal{D}} \cup \{(d, (v, 0)) \mid v \in \eta(d)\} \cup \{(v, i), (v', i)\} \mid (v, v') \in E_{\mathcal{F}}, i \in \omega\} \cup \{(v, i), (v', i+1)\} \mid v' \in \sigma_i(v), i \in \omega\}$.

Thus G_{α} is obtained by taking \mathcal{D} together with an infinite disjoint union of \mathcal{F} such that edges between \mathcal{D} and the first copy of \mathcal{F} are put according to the mapping η , and edges between successive copies of \mathcal{F} are put according to σ_i .

Figure 3 illustrates the general shape of a unary automatic graph of finite degree that is build from \mathcal{D} , \mathcal{F} , η , and σ^{ω} , where σ^{ω} is the infinite word $\sigma\sigma\sigma \dots$.




FIGURE 3. Unary automatic graph of finite degree $G_{\eta\sigma^{\omega}}$

Theorem 4.4. *A graph of finite degree $\mathcal{G} = (V, E)$ possesses a unary automatic presentation if and only if there exist finite graphs \mathcal{D}, \mathcal{F} and mappings $\eta : V_{\mathcal{D}} \rightarrow P(V_{\mathcal{F}})$ and $\sigma : V_{\mathcal{F}} \rightarrow P(V_{\mathcal{F}})$ such that \mathcal{G} is isomorphic to $\mathcal{G}_{\eta\sigma^{\omega}}$.*

Proof. Let $G = (V, E)$ be a unary automatic graph of finite degree. Let \mathcal{A} be an automaton recognizing E . In linear time on the number of states of \mathcal{A} we can easily transform \mathcal{A} into a one-loop automaton. So, we assume that \mathcal{A} is a one-loop automaton with loop constant p . We construct the finite graph \mathcal{D} by setting $V_{\mathcal{D}} = \{q_0, q_1, \dots, q_{p-1}\}$, where q_0 is the starting state, q_0, \dots, q_{p-1} are all states on the $(1, 1)$ -tail such that q_i is reached from q_{i-1} by reading $(1, 1)$ for $i > 0$; and for $0 \leq i \leq j < p$, $(q_i, q_j) \in E_{\mathcal{D}}$ iff there is a final state q_f on the $(\diamond, 1)$ -tail out of q_i , and the distance from q_i to q_f is $j - i$. We construct the graph \mathcal{F} similarly by setting $V_{\mathcal{F}} = \{q'_0, \dots, q'_{p-1}\}$ where q'_0, \dots, q'_{p-1} are all states on the $(1, 1)$ -loop. The edge relation $E_{\mathcal{F}}$ is defined in a similar way as $E_{\mathcal{D}}$. The mapping $\eta : V_{\mathcal{D}} \rightarrow P(V_{\mathcal{F}})$ is defined for any $m, n \in \{0, \dots, p-1\}$ by putting q'_n in $\eta(q_m)$ if and only if there exists a final state q_f on the $(\diamond, 1)$ -tail out of q_m , and the distance from q_m to q_f equals $p + n - m$. The mapping σ is constructed in a similar manner by reading the $(\diamond, 1)$ -tails out of the $(1, 1)$ -loop. It is clear from this construction that the graphs G and $G_{\eta\sigma^{\omega}}$ are isomorphic.

Conversely, consider the graph $G_{\eta\sigma\omega}$ for some $\eta \in \Sigma_{\mathcal{D}}$ and $\sigma \in \Sigma_{\mathcal{F}}$. Assume that $V_{\mathcal{D}} = \{q_0, \dots, q_{\ell-1}\}$, $V_{\mathcal{F}} = \{q'_0, \dots, q'_{p-1}\}$. A one-loop automaton \mathcal{A} recognizing the edge relation of $G_{\eta\sigma\omega}$ is constructed as follows. The $(1, 1)$ -tail of the automaton is formed by $\{q_0, \dots, q_{\ell-1}\}$ and the $(1, 1)$ -loop is formed by $\{q'_0, \dots, q'_{p-1}\}$, both in natural order. The initial state is q_0 . If for some $i < j$, $\{q_i, q_j\} \in E_{\mathcal{D}}$, then put a final state q_f on the $(\diamond, 1)$ -tail starting from q_i such that the distance from q_i to q_f is $j - i$. If $q'_j \in \eta(q_i)$, then repeat the process but make the corresponding distance $p + j - i$. The set of edges $E_{\mathcal{F}}$ and mapping σ are treated in a similar manner by putting final states on the $(\diamond, 1)$ -tails from the $(1, 1)$ -loop.

Again, we see that \mathcal{A} represents a unary automatic graph that is isomorphic to $G_{\eta\sigma\omega}$. \square

If \mathcal{G} is a unary automatic graph of finite degree, the parameters \mathcal{D} , \mathcal{F} , σ and η can be extracted in $O(p^2)$ time, where p is the loop constant of the one-loop automaton representing the graph. Furthermore, $|V_{\mathcal{F}}| = |V_{\mathcal{D}}| = p$.

5. DECIDING THE INFINITE COMPONENT PROBLEM

A **component** of G is the transitive closure of a vertex under the edge relation. The **infinite component problem** asks whether a given graph \mathcal{G} has an infinite component.

Theorem 5.1. *The infinite component problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(p^3)$, where p is the loop constant of the unary automaton recognizing \mathcal{G} .*

By Theorem 4.4, let $G = G_{\eta\sigma\omega}$. We observe that it is sufficient to consider the case in which $\mathcal{D} = \emptyset$ (hence $G = G_{\sigma\omega}$) since $G_{\eta\sigma\omega}$ has an infinite component if and only if $G_{\sigma\omega}$ has one.

Let \mathcal{F}^i be the i^{th} copy of \mathcal{F} in \mathcal{G} . Let x^i be the copy of vertex x in \mathcal{F}^i . We construct a finite directed graph $\mathcal{F}^\sigma = (V^\sigma, E^\sigma)$ as follows. Each node in V^σ represents a distinct connected component in \mathcal{F} . For simplicity, we assume that $|V^\sigma| = |V_{\mathcal{F}}|$ and hence use x to denote its own component in \mathcal{F} . The case in which $|V^\sigma| < |V_{\mathcal{F}}|$ can be treated in a similar way. For $x, y \in V_{\mathcal{F}}$, put $(x, y) \in E^\sigma$ if and only if $y' \in \sigma(x')$ for some x' and y' that are in the same component as x and y , respectively. Constructing \mathcal{F}^σ requires finding connected components of \mathcal{F} hence takes time $O(p^2)$. To prove the above theorem, we make essential use of the following definition which is taken from [7].

Definition 5.2. *An **oriented walk** in a directed graph G is a subgraph \mathcal{P} of G that consists of a sequence of nodes v_0, \dots, v_k such that for $1 \leq i \leq k$, either (v_{i-1}, v_i) or (v_i, v_{i-1}) is an arc in G , and for each $0 \leq i < k$, exactly one of (v_{i-1}, v_i) and (v_i, v_{i-1}) belongs to \mathcal{P} . An oriented walk is an **oriented cycle** if $v_0 = v_k$ and there are no repeated nodes in v_1, \dots, v_k .*

In an oriented walk \mathcal{P} , an arc (v_i, v_{i+1}) is called a **forward arc** and (v_{i+1}, v_i) is called a **backward arc**. The **net length** of \mathcal{P} , denoted $\text{disp}(\mathcal{P})$, is the difference between the number of forward arcs and backward arcs. Note the net length can be negative. Given an oriented walk $\mathcal{P} = v_0, \dots, v_m$, we define the **low point** of \mathcal{P} as $\min\{\text{disp}(v_0 \dots v_\ell) \mid 1 \leq \ell \leq m\}$. The low point of oriented walk \mathcal{P} is at most $\min\{0, \text{disp}(\mathcal{P})\}$, and hence is not positive. The next lemma establishes a connection between oriented walks in \mathcal{F}^σ and paths in G .

Lemma 5.3. *Let \mathcal{P} be an oriented walk from x to y whose net length is d and low point is $-\ell$. For every $i \geq \ell$, the oriented walk \mathcal{P} defines a path P^i in \mathcal{G} from x^i to y^{i+d} . Moreover, the smallest j such that $P^i \cap \mathcal{F}^j \neq \emptyset$ is equal to $i - \ell$.*

Proof. The path P^i is obtained by starting at vertex x^i in G and following the edges in G determined by the sequence of edges in \mathcal{P} . \square

Lemma 5.4. *There is an infinite component in G if and only if there is an oriented cycle in \mathcal{F}^σ such that the net length of the cycle is positive.*

Proof. Suppose there is an oriented cycle \mathcal{P} from x to x in \mathcal{F}^σ of net length $m > 0$ and low point $-\ell$. By the previous lemma, for all $i > \ell$, \mathcal{P} defines the path P_i in G from x^i to x^{i+m} where P_i lies in $\mathcal{F}^{i-\ell} \cup \dots \cup \mathcal{F}^{i+m}$. Therefore, for a fixed $i > \ell$, all vertices in the set $\{x^{j+m+i} \mid j \in \omega\}$ belong to the same component of G . In particular, this implies that G contains an infinite component.

Conversely, suppose there is an infinite component D in G . Since \mathcal{F} is finite, there must be some x in $V_{\mathcal{F}}$ such that there are infinitely many copies of x in D . Let x^i and x^j be two copies of x in D such that $i < j$. Consider a path between x^i and x^j . We can assume that on this path there is at most one copy of any vertex $y \in V_{\mathcal{F}}$ apart from x (otherwise, choose x^j to be the copy of x in the path that has this property). By definition of $G_{\sigma\omega}$ and \mathcal{F}^σ , the node x must be on an oriented cycle of \mathcal{F}^σ with net length $j - i$. \square

Proof of Theorem 5.1. By the equivalence in Lemma 5.4, it suffices to provide an algorithm that decides if \mathcal{F}^σ contains an oriented cycle with positive net length. Notice that the existence of an oriented cycle with positive net length is equivalent to the existence of an oriented cycle with negative net length. Therefore, we give an algorithm which finds oriented cycles with non-zero net length.

For each node x in \mathcal{F}^σ , we search for an oriented cycle of positive net length from x by creating a labelled queue of nodes Q_x which are connected to x .

ALG:Oriented-Cycle

- (1) Pick node $x \in \mathcal{F}^\sigma$ for which a queue has not been built yet. Initially the queue Q_x is empty. Let $d(x) = 0$, and put x into the queue. Mark x as *unprocessed*. If queues have been built for each $x \in \mathcal{F}^\sigma$, stop the process and return *NO*.
- (2) Let y be the first *unprocessed* node in Q_x . If there are no *unprocessed* nodes in Q_x , return to (1).
- (3) For each of the nodes z in the set $\{z \mid (y, z) \in E^\sigma \text{ or } (z, y) \in E^\sigma\}$, do the following.
 - (a) If $(y, z) \in E^\sigma$, set $d'(z) = d(y) + 1$; if $(z, y) \in E^\sigma$, set $d'(z) = d(y) - 1$. (If both hold, do steps (a), (b), (c) first for (z, y) and then for (y, z) .)
 - (b) If $z \notin Q_x$, then set $d(z) = d'(z)$, put z into Q_x , and mark z as *unprocessed*.
 - (c) If $z \in Q_x$ then
 - (i) if $d(z) = d'(z)$, move to next z ,
 - (ii) if $d(z) \neq d'(z)$, stop the process and return *YES*.
- (4) Mark y as *processed* and go back to (2).

An important property of this algorithm is that when we are building a queue for node x and are processing z , both $d(z)$ and $d'(z)$ represent net lengths of paths from x to z .

We claim that the algorithm returns *YES* if and only if there is an oriented cycle in \mathcal{F}^σ with non-zero net length. Suppose the algorithm returns *YES*. Then, there is a base node x and a node z such that $d(z) \neq d'(z)$. This means that there is an oriented walk \mathcal{P} from x to z with net length $d(z)$ and there is an oriented walk $\overleftarrow{\mathcal{P}'}$ from x to z with net length $d'(z)$.

Consider the oriented walk $\mathcal{P}\overleftarrow{\mathcal{P}'}$, where $\overleftarrow{\mathcal{P}'}$ is the oriented walk \mathcal{P}' in reverse direction. Clearly this is an oriented walk from x to x with net length $d(z) - d'(z) \neq 0$. If there is no repeated nodes in $\overleftarrow{\mathcal{P}'}$, then it is the required oriented cycle. Otherwise, let y be a repeated node in $\overleftarrow{\mathcal{P}'}$ such that no nodes between the two occurrences of y are repeated. Consider the oriented walk between these two occurrences of y , if it has a non-zero net length, then it is our required oriented cycle; otherwise, we disregard the part between the two occurrences of z and make the oriented walk shorter without altering its net length.

Conversely, suppose there is an oriented cycle $\mathcal{P} = x_0, \dots, x_m$ of non-zero net length where $x_0 = x_m$. However, we assume for a contradiction that the algorithm returns *NO*. Consider how the algorithm acts when we pick x_0 at step (1). For each $0 \leq i \leq m$, one can prove the following statements by induction on i .

(\star) x_i always gets a label $d(x_i)$

($\star\star$) $d(x_i)$ equals the net length of the oriented walk from x_0 to x_i in \mathcal{P} .

By the description of the algorithm, x_0 gets the label $d(x_0) = 0$. Suppose the statements holds for x_i , $0 \leq i < m$, then at the next stage, the algorithm labels all nodes in $\{z \mid (z, x_i) \in E^\sigma \text{ or } (x_i, z) \in E^\sigma\}$. In particular, it calculates $d'(x_{i+1})$. By the inductive hypothesis, $d'(x_{i+1})$ is the net length of the oriented walk from x_0 to x_{i+1} in \mathcal{P} . If x_{i+1} has already had a label $d(x_{i+1})$ and $d(x_{i+1}) \neq d'(x_{i+1})$, then the algorithm would return *YES*. Therefore $d(x_{i+1}) = d'(x_{i+1})$. By assumption on \mathcal{P} , $d(x_m) \neq 0$. However, since $x_0 = x_m$, the induction gives that $d(x_m) = d(x_0) = 0$. This is a contradiction, and thus the above algorithm is correct.

In summary, the following algorithm solves the infinite component problem. Suppose we are given a unary automaton (with loop constant p) which recognizes the unary automatic graph of finite degree G . Recall that p is also the cardinality of $V_{\mathcal{F}}$. We first compute \mathcal{F}^σ , in time $O(p^2)$. Then we run `Edge-Cycle` to decide whether \mathcal{F}^σ contains an oriented cycle with positive net length. For each node x in \mathcal{F}^σ , the process runs in time $O(p^2)$. Since \mathcal{F}^σ contains p number of nodes, this takes time $O(p^3)$.

Note that Lemma 5.4 holds for the case when $|V_{\mathcal{F}}| > |\mathcal{F} / \sim_{comp}|$. Therefore the algorithm above can be slightly modified to apply to this case as well. \square

6. DECIDING INFINITY TESTING PROBLEM

We next turn our attention to the **infinity testing problem** for unary automatic graphs of finite degree. Recall that this problem asks for an algorithm that, given a vertex v and a graph G , decides if v belongs to an infinite component. We prove the following theorem.

Theorem 6.1. *The infinity testing problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(p^5)$, where p is the loop constant of the unary automaton \mathcal{A} recognizing \mathcal{G} . In particular, when \mathcal{A} is fixed, there is a constant time algorithm that decides the infinity testing problem on \mathcal{G} .*

We first need the following lemma. Recall the notation

Lemma 6.2. *Let $x \in V_{\mathcal{F}}$. If x^i belongs to an infinite component of G then for all $j > 0$, x^{i+j} also belongs to an infinite component of G .*

Proof. Suppose x^i belongs to an infinite component. Take any y^k that is in the same component as x^i with $k \geq i$. Consider a path P in the graph G between x^i and y^k . Note that for any edge (w^s, z^t) in the path P , there must also be the edge (w^{s+j}, z^{t+j}) in G for all $j > 0$. Hence, in G there is a path between x^{i+j} and y^{k+j} . This proves that x^{i+j} also belongs to an infinite component of G . \square

In the previous section, we saw a connection between infinite components in G and oriented cycles in \mathcal{F}^σ . We now explore this connection further. Let \mathcal{C} be the set of nodes x in \mathcal{F}^σ for which there exists an oriented cycle from x with positive net length and low point 0. We have the following lemma.

Lemma 6.3. *If $x \in \mathcal{C}$, then x^i is in an infinite component for all $i \in \omega$.*

Proof. If $x \in \mathcal{C}$, then by the proof of Lemma 5.4, x^0 must be in an infinite component. By Lemma 6.2, x^i is in an infinite component for all $i \in \omega$. \square

We call an oriented walk **simple** if the oriented walk doesn't contain any repeated nodes. For any $k \geq 0$, let $\mathcal{C}[k]$ be the set of all nodes $x \notin \mathcal{C}[0] \cup \dots \cup \mathcal{C}[k-1]$ that can reach \mathcal{C} by a simple oriented walk with low point $-k$. Note that this definition ensures that $\mathcal{C} \subseteq \mathcal{C}[0]$. Moreover, since $|\mathcal{F}^\sigma| \leq p$, a simple oriented walk may have at most p steps and hence $\mathcal{C}[k] = \emptyset$ for $k > p-1$.

Lemma 6.4. *For each vertex x^i , x^i belongs to an infinite component in G if and only if node $x \in \mathcal{C}[k]$ for some $0 \leq k \leq \min\{i, p-1\}$.*

Proof. Suppose $x \in \mathcal{C}[k]$ for some $0 \leq k \leq i$. Take the oriented walk from x to a node $y \in \mathcal{C}$ with low point $-k$ as stipulated by definition of $\mathcal{C}[k]$. There is a path in G from x^i to y^j where $j \geq i-k$. By Lemma 6.3, y^j is in an infinite component. Therefore x^i is in an infinite component.

It remains to prove that if x^i belongs to an infinite component, then there exists an oriented walk from x to \mathcal{C} with low point $-k$, where $0 \leq k \leq \min\{i, p-1\}$.

Let D be the infinite component of x_i . Since \mathcal{F} is finite, there must be some y in $V_{\mathcal{F}}$ such that D contains infinitely many copies of y . Let y^s and y^t be two copies of y in D with $s < t$. Take a path P in G between y^s and y^t such that P contains no more than one copy of each vertex in $V_{\mathcal{F}}$ apart from y . Let m be the least number such that $P \cap \mathcal{F}^m \neq \emptyset$. Take a vertex z^m in the path P . Then P is divided into two paths P_1 and P_2 where P_1 goes from y^s to z^m , and P_2 goes from z^m to y^t . Hence there must also be a path P_3 from y^t to z^{m+t-s} . By joining P_2 and P_3 together we obtain a path between z^m and z^{m+t-s} . This path defines an oriented cycle in \mathcal{F}^σ with net length > 0 and low point 0. Hence, $z \in \mathcal{C}$.

Take a path in G between x^i and a copy of z in D such that no more than one copy of each vertex in \mathcal{F} appears in the path. This path defines an oriented walk in \mathcal{F}^σ from x to z with low point $\leq \min\{i, p-1\}$. \square

Lemma 6.5. *Let \mathcal{G} be a unary automatic graph of finite degree presented by a unary automaton \mathcal{A} with loop constant p . There is an algorithm to compute \mathcal{C} for \mathcal{G} in time $O(p^4)$.*

Proof. The algorithm first computes the directed graph \mathcal{F}^σ in time $O(p^2)$. Next, for each $x \in \mathcal{F}^\sigma$, we perform a breadth-first search through \mathcal{F}^σ for oriented walks starting at x . In order to compose the path \mathcal{P} , we put (y, d) in the queue, where y is the incremental destination of \mathcal{P} and d is its net length. We then compute the following properties of the pair (y, d) :

- (1) $level(y, d)$ is the length of the oriented walk \mathcal{P} from x to y ; and
- (2) $path(y, d)$ is a sequence of pairs $(x_0, d_0) \dots (x_{level(y,d)}, d_{level(y,d)})$ representing the initial segment of the oriented walk \mathcal{P} . Note that $(x_0, d_0) = (x, 0)$ and that $(x_{level(y,d)}, d_{level(y,d)}) = (y, d)$.

More precisely, the algorithm for a given $x \in \mathcal{F}^\sigma$ is as follows.

ALG:C-Membership

- (1) Initially the queue Q is empty. Put $(x, 0)$ into the queue. Mark $(x, 0)$ as *unprocessed* and set $level(x, 0) = 0$, $path(x, 0) = (x, 0)$.
- (2) If no *unprocessed* pair is left in the queue, stop and output *NO*. Otherwise, take the first *unprocessed* (y, d) in Q .
- (3) If $level(y, d) \geq p$, then stop and output *NO*.
- (4) For arcs e of the form (y, z) or (z, y) in E^σ do the following:
 - (a) If $e = (y, z)$, set $j = d + 1$; if $e = (z, y)$, set $j = d - 1$.
 - (b) If $z = x$ and $j > 0$, stop the process and return *YES*.
 - (c) If (z, d') is not in $path(y, d)$ for any d' , and if $j \geq 0$ and $(z, j) \notin Q$, then we wish to process (z, j) :
 - (i) Put (z, j) into Q and mark (z, j) as *unprocessed*.
 - (ii) Set $level(z, j) = level(y, d) + 1$ and set $path(z, j) = path(y, d) \cdot (z, j)$.
- (5) Mark (y, d) as *processed* and go back to (2).

We claim that running **C-Membership** with input x returns *YES* if and only if $x \in \mathcal{C}$. Suppose that the algorithm returns *YES*. Then there exists a simple oriented walk \mathcal{P} from x to x with positive net length. Let \mathcal{P} be x_0, \dots, x_m such that $x_0 = x_m = x$. The algorithm ensures that the

net length of the (sub)oriented walk in \mathcal{P} from x_0 to each x_i is non-negative. Thus the low point of \mathcal{P} is greater or equal to 0. Hence, $x \in \mathcal{C}$.

For the other direction, suppose there exists an oriented cycle $\mathcal{P} = x_0, \dots, x_m$ of positive net length and zero low point. Suppose the algorithm does not return *YES*. Run the algorithm from x_0 . For all x_i , one can prove the following statements by induction.

(\star) There exists $d_i \geq 0$ such that $(x_i, d_i) \in Q$.

($\star\star$) d_i equals the net length of the oriented walk from x_0 to x_i in \mathcal{P} .

Note that $level(y, d)$ of any $(y, d) \in Q$ must be less than p . Since every time the level is incremented by 1 the net length is either incremented by 1 or decremented by 1, hence d must also be less than p . Therefore, the number of elements in Q is bounded above by p^2 and so for each $x \in \mathcal{F}^\sigma$, the algorithm takes time $O(p^3)$. To compute \mathcal{C} , we need to run **C-Membership** on every x in \mathcal{F}^σ , putting x in \mathcal{C} if and only if the algorithm returns *YES*. This takes time $O(p^4)$. \square

Using Lemma 6.5, we can iteratively compute the set $\mathcal{C}[k]$ for any $0 \leq k \leq p-1$ as follows. For each node $x \notin \mathcal{C}[0] \cup \dots \cup \mathcal{C}[k-1]$, we run operations similar to the ones described above. The differences are the following:

(1) At step (4)(b), the process stops and returns *YES* whenever $z \in \mathcal{C}$ and $j \geq -k$.

(2) At step (4)(c), the process puts a pair (z, j) into the queue Q if $j \geq -k$ and $(z, j) \notin Q$.

The proof of correctness can be done in a similar manner as the proof above. This algorithm also takes times $O(p^4)$.

Proof of Theorem 6.1. We assume the input vertex x^i is given by tuple (x, i) . By Lemma 6.4, to check if x^i is in an infinite component, the algorithm needs to compute $\mathcal{C}[0], \dots, \mathcal{C}[\min\{i, p-1\}]$. As a consequence of Lemma 6.5, this takes time $O(p^5)$. The algorithm then checks whether $x \in \mathcal{C}[k]$ for some $0 \leq k \leq \min\{i, p-1\}$. Since x is bounded by p , if \mathcal{A} is fixed, once the sets $\mathcal{C}[0], \dots, \mathcal{C}[p-1]$ are found, checking whether x^i belongs to an infinite component takes constant time. \square

7. DECIDING REACHABILITY PROBLEM

Suppose G is a unary automatic graph of finite degree represented by a unary automaton with loop constant p . The **reachability problem** on G is formulated as: given two vertices x^i, y^j in G , decide if x^i and y^j are in the same component. We prove the following theorem.

Theorem 7.1. *Suppose \mathcal{G} is a unary automatic graph of finite degree represented by unary automaton \mathcal{A} of loop constant p . There exists a polynomial-time algorithm that solves the reachability problem on \mathcal{G} . For inputs u, v , the running time of the algorithm is $O(|u| + |v| + p^5)$.*

We restrict to the case when $G = G_{\sigma^\omega}$. The proof can be modified slightly to work in the more general case, $G = G_{\eta\sigma^\omega}$.

Since, by Theorem 6.1, there is an $O(p^5)$ -time algorithm to check if x^i is in a finite component, we can work on the two possible cases separately. We first deal with the case when the input x^i is in a finite component. We have the following lemma.

Lemma 7.2. *If x^i is in a finite component, then x^i and y^j are in the same component only if $i - p < j < i + p$.*

Proof. Suppose the contrary. Then y^j is in the same component as x^i and we can assume without loss of generality that $j \geq i + p$ (the case where $i - p \geq j$ can be treated similarly). Then take a path P in G from x^i to y^j . Since p is the cardinality of $V_{\mathcal{F}}$, there is $z \in V_{\mathcal{F}}$ such that z^s and z^t appear in P with $s < t$. Therefore all nodes in the set $\{z^{s+(t-s)m} \mid m \in \omega\}$ are in the same component as x^i , contradicting the finiteness of the component of x^i . \square

Suppose x^i is in a finite component, let $i' = \min\{p, i\}$. To decide if x^i and y^j are in the same component, we run a breadth first search in G starting from $x^{i'}$ and going through all vertices in $\mathcal{F}^{i-i'}, \dots, \mathcal{F}^{i+p}$. This is sufficient by the lemma above. The algorithm is described as follows.

ALG: FiniteReach

- (1) Let $i' = \min\{p, i\}$.
- (2) Initialize the queue Q to be empty. Put the pair $(x, 0)$ into Q and mark it as *unprocessed*.
- (3) If there are no *unprocessed* pairs in Q , stop the process. Otherwise, let (y, d) be the first *unprocessed* pair. For arcs e of the form (y, z) or (z, y) in E^σ , do the following.
 - (a) If e is of the form (y, z) , let $d' = d + 1$; if e is of the form (z, y) , let $d' = d - 1$.
 - (b) If $-i' \leq d' \leq p$ and (z, d') is not in Q , then put (z, d') into Q and mark (z, d') as *unprocessed*.
- (4) Mark (y, d) as *processed*, and go to (2).

By Lemma 7.2, x^i and y^j are in the same (finite) component if and only if after running **FiniteReach** on the input x^i , the pair $(y, j - i)$ is in Q . Note that the running time of **FiniteReach** is bounded by the number of edges in G restricted to $\mathcal{F}^0, \dots, \mathcal{F}^{2p}$. Therefore the running time is $O(p^3)$.

Corollary 7.3. *If all components of G are finite and we represent (x^i, y^j) as $(x^i, y^j, j - i)$, then there is an $O(p^3)$ -algorithm deciding if x^i and y^j are in the same component. \square*

Now, suppose that x^i is in an infinite component. We start with the following question: given $y \in V_{\mathcal{F}}$, are x^i and y^i in the same component in G ? To answer this, we present an algorithm that computes all nodes $y \in V_{\mathcal{F}}$ whose i^{th} copy lies in the same G -component as x^i . The algorithm is similar to **FiniteReach**, except that it does not depend on the input i . Line(3b) in the algorithm is changed to the following:

- (3b) If $-p \leq d' \leq p$ and (z, d') is not in Q , then put (z, d') into Q and mark (z, d') as *unprocessed*.

We use this modified algorithm to define the set $Reach(x) = \{y \mid (y, 0) \in Q\}$. Intuitively, we can think of the algorithm as a breadth first search through $\mathcal{F}^0 \cup \dots \cup \mathcal{F}^{2p}$ which originates at x^p . Therefore, $y \in Reach(x)$ if and only if there exists a path from x^p to y^p in G restricted to $\mathcal{F}^0 \cup \dots \cup \mathcal{F}^{2p}$.

Lemma 7.4. *Suppose x^i is in an infinite component. The node y^i is in the same component as x^i if and only if y^i is also in an infinite component and $y \in Reach(x)$.*

Proof. Suppose y^i is in an infinite component and $y \in Reach(x)$. If $i \geq p$, then the observation above implies that there is a path from x^i to y^i in $\mathcal{F}^{i-p} \cup \dots \cup \mathcal{F}^{i+p}$. So, it remains to prove that x^i and y^i are in the same component even if $i < p$.

Since $y \in Reach(x)$, there is a path P in G from x^p to y^p . Let ℓ be the least number such that $\mathcal{F}^\ell \cap P \neq \emptyset$. If $i \geq p - \ell$, then it is clear that x^i and y^i are in the same component. Thus, suppose that $i < p - \ell$. Let z be a node such that $z^\ell \in P$. Then P is P_1P_2 where P_1 is a path from x^p to z^ℓ and P_2 is a path from z^ℓ to y^p . By Lemma 6.2, since x^i is in an infinite component, so is x^p . There exists an $r > 0$ such that all nodes in the set $\{x^{p+rm} \mid m \in \omega\}$ are in the same component. Likewise, there is an $r' > 0$ such that all nodes in $\{y^{p+r'm} \mid m \in \omega\}$ are in the same component. Consider $x^{p+rr'}$ and $y^{p+rr'}$. Analogous to the path P_1 , there is a path P'_1 from $x^{p+rr'}$ to $z^{\ell+rr'}$. Similarly, there is a path P'_2 from $z^{\ell+rr'}$ to $y^{p+rr'}$. We describe another path P' from x^p to y^p as follows. P' first goes from x^p to $x^{p+rr'}$, then goes along $P'_1P'_2$ from $x^{p+rr'}$ to $y^{p+rr'}$ and finally goes to y^p . Notice that the least ℓ' such that $\mathcal{F}_{\ell'} \cap P' \neq \emptyset$ must be larger than ℓ . We can iterate this procedure of lengthening the path between x^p and y^p until $i < p - \ell'$, as is required to reduce to the previous case.

To prove the implication in the other direction, we assume that x^i and y^i are in the same infinite component. Then y^i is, of course, in an infinite component. We want to prove that $y \in Reach(x)$.

Let $i' = \min\{p, i\}$. Suppose there exists a path P in G from x^i to y^i which stays in $\mathcal{F}^{i-i'} \cup \dots \cup \mathcal{F}^{i+p}$. Then, indeed, $y \in \text{Reach}(x)$. On the other hand, suppose no such path exists. Since x^i and y^i are in the same component, there is some path P from x^i to y^i . Let $\ell(P)$ be the largest number such that $P \cap \mathcal{F}^{\ell(P)} \neq \emptyset$. Let $\ell'(P)$ be the least number such that $P \cap \mathcal{F}^{\ell'(P)} \neq \emptyset$. We are in one of two cases: $\ell(P) > i+p$ or $\ell'(P) < i-p$. We will prove that if $\ell(P) > i+p$ then there is a path P' from x^i to y^i such that $\ell(P') < \ell(P)$ and $\ell'(P') \geq i-p$. The case in which $\ell'(P) < i-p$ can be handled in a similar manner.

Without loss of generality, we assume $\ell'(P) = i$ since otherwise we can change the input x and make $\ell'(P) = i$. Let z be a node in \mathcal{F} such that $z^{\ell(P)} \in P$. Then P is P_1P_2 where P_1 is a path from x^i to $z^{\ell(P)}$ and P_2 is a path from $z^{\ell(P)}$ to y^i . Since $\ell(P) > i+p$, there must be some s^j and s^{j+k} in P_1 such that $k > 0$. For the same reason, there must be some t^m and t^{m+n} in P_2 such that $n > 0$. Therefore, P contains paths between any consecutive pair of nodes in the sequence $(x^i, s^j, s^{k+j}, z^p, t^{m+n}, t^n, y^i)$. Consider the following sequence of nodes:

$$(x^i, s^j, t^{m+n-k}, t^{n-k}, s^{j-m}, s^{j+k-m}, t^n, y^i).$$

It is easy to check that there exists a path between each pair of consecutive nodes in the sequence. Therefore the above sequence describes a path P' from x^i to y^i . It is easy to see that $\ell(P') = \ell(P) - n$. Also since $\ell'(P) = i$, $\ell'(P') > i-p$. Therefore P' is our desired path. \square

In the following, we abuse notation by using Reach and σ on subsets of $V_{\mathcal{F}}$. We inductively define a sequence $Cl_0(x), Cl_1(x), \dots$ such that each $Cl_k(x)$ is a subset of $V_{\mathcal{F}}$. Let $Cl_0(x) = \text{Reach}(x)$ and for $k > 0$, we define $Cl_k(x) = \text{Reach}(\sigma(Cl_{k-1}(x)))$. The following lemma is immediate from this definition.

Lemma 7.5. *Suppose x^i is in an infinite component, then x^i and y^j are in the same component if and only if y^j is also in an infinite component and $y \in Cl_{j-i}(x)$.* \square

We can use the above lemma to construct a simple-minded algorithm that solves the reachability problem on inputs x^i, y^j .

ALG: NaïveReach

- (1) Check if each of x^i, y^j are in an infinite component of G (using the algorithm of Theorem 6.1).
- (2) If exactly one of x^i and y^j is in a finite component, then return *NO*.
- (3) If both x^i and y^j are in finite components, then run **FiniteReach** on input x^i and check if $(y, j-i)$ is in Q .
- (4) If both x^i and y^j are in infinite components, then compute $Cl_{j-i}(x)$. If $y \in Cl_{j-i}(x)$, return *YES*; otherwise, return *NO*.

We now consider the complexity of this algorithm. The set $Cl_0(x)$ can be computed in time $O(p^3)$. Given $Cl_{k-1}(x)$, we can compute $Cl_k(x)$ in time $O(p^4)$. Therefore, the total running time of **NaïveReach** on input x^i, y^j is $(j-i) \cdot p^4$. We want to replace the multiplication with addition and hence tweak the algorithm.

From Lemma 6.4, x^i is in an infinite component in G if and only if there is an oriented cycle \mathcal{C} with positive net length and zero low point and which is reachable from x via a simple oriented walk with low point $\geq -i$. Now, suppose that x^i is in an infinite component. We can use the algorithm that solves the infinity testing problem to find such an oriented cycle \mathcal{C} . Let r be the net length of \mathcal{C} . Note that r can be computed from the algorithm. It is easy to see that all nodes in the set $\{x^{i+mr} \mid m \in \omega\}$ belong to the same component.

Lemma 7.6. $Cl_0(x) = Cl_r(x)$.

Proof. By definition, $y \in Cl_0(x)$ if and only if x^p and y^p are in the same component of G . Suppose that there exists a path in G from x^p to y^p . Then there is a path from x^{p+r} to y^{p+r} . Since x^p and x^{p+r} are in the same component of G , x^p and y^{p+r} are in the same component. Hence $y \in Cl_r(x)$.

For the reverse inclusion, suppose $y \in Cl_r(x)$. Then there exists a path from x^p to y^{p+r} . Therefore, x^{p+r} and y^{p+r} are in the same component. Since $r \leq p$, x^p and y^p are in the same component. \square

Using the above lemma, we define a new algorithm **Reach** on inputs x^i, y^j by replacing line (4) in **NaiveReach** with

(4) If x^i and y^j belong to infinite components, then compute $Cl_0(x), \dots, Cl_{r-1}(x)$. If $y \in Cl_k(x)$ for $k < r$ such that $j - i = k \pmod r$, return *YES*; otherwise, return *NO*.

Proof of Theorem 7.1. Say input vertices are given as x^i and y^j . By Lemma 7.5 and Lemma 7.6, the algorithm **Reach** returns *YES* if and only if x^i and y^j are in the same component. Since $r \leq p$, calculating $Cl_0(x), \dots, Cl_{r-1}(x)$ requires time $O(p^5)$. Therefore the running time of **Reach** on input x^i, y^j is $O(i + j + p^5)$. \square

Notice that, in fact, the algorithm produces a number $k < p$ such that in order to check if x^i, y^j ($j > i$) are in the same component, we need to test if $j - i < p$ and if $j - i = k \pmod p$. Therefore if \mathcal{G} is fixed, we compute $Cl_0(x), \dots, Cl_{r_x-1}(x)$ for all x beforehand, then deciding whether two vertices u, v belong to the same component takes linear time. The above proof can also be used to build a unary automaton that decides reachability uniformly:

Corollary 7.7. *Given a unary automatic graph of finite degree \mathcal{G} represented by a unary automaton with loop constant p , there is a deterministic automaton with at most $2p^4 + p^3$ states that solves the reachability problem on \mathcal{G} . The time required to construct this automaton is $O(p^6)$.*

Proof. For all $0 \leq x < p, i \in \omega$, let string 1^{ip+x} represent vertex x^i in G . Suppose $ip + x \leq jp + y$, we construct a unary automaton \mathcal{A}_{Reach} that accepts $(1^{ip+x}, 1^{jp+y})$ if and only if x^i and y^j are in the same component in G .

- (1) \mathcal{A}_{Reach} has a $(1, 1)$ -tail of length p^2 . Let the states on the tail be $q_0, q_1, \dots, q_{p^2-1}$, where q_0 is the initial state. These states represent nodes in $\mathcal{F}^0, \mathcal{F}^1, \dots, \mathcal{F}^{p-1}$.
- (2) From q_{p^2-1} , there is a $(1, 1)$ -loop of length p . We call the states on the loop be $q'_0, q'_1, \dots, q'_{p-1}$. These states represent nodes in \mathcal{F}^p .
- (3) For $0 \leq x, i < p$, there is a $(\diamond, 1)$ -tail from q_{ip+x} of length $p^2 - x$. We denote the states on this tail by $q_{ip+x}^1, \dots, q_{ip+x}^{p^2-x}$. These states represent nodes in $\mathcal{F}^i, \mathcal{F}^{i+1}, \dots, \mathcal{F}^{i+p-1}$.
- (4) For $0 \leq x, i \leq p$, if x^i is in an infinite component, then there is a $(\diamond, 1)$ -loop of length $r \times p$ from $q_{ip+x}^{p^2-x}$. The states on this loop are called $\tilde{q}_{ip+x}^1, \dots, \tilde{q}_{ip+x}^{rp}$. These states represent nodes in $\mathcal{F}^{i+p}, \dots, \mathcal{F}^{i+p+r-1}$.
- (5) For $0 \leq x \leq p$, if x^p is in a finite component, then there is a $(\diamond, 1)$ -tail from q'_x of length p^2 . These states are denoted $\hat{q}_x^1, \dots, \hat{q}_x^{p^2}$ and represent nodes in $\mathcal{F}_p, \dots, \mathcal{F}_{2p-1}$.
- (6) If x^p is in an infinite component, from q'_x , there is a $(\diamond, 1)$ -loop of length $r \times p$. We write these states as $\tilde{q}_x^1, \dots, \tilde{q}_x^{rp}$.

The final (accepting) states of \mathcal{A}_{Reach} are defined as follows:

- (1) States $q_0, \dots, q_{p^2-1}, q'_0, \dots, q'_{p-1}$ are final.
- (2) For $i < p$, if x^i is in a finite component, run the algorithm **FiniteReach** on input x^i and declare state q_{ip+x}^{jp+y-x} final if $(y, j) \in Q$.
- (3) For $i < p$, if x^i is in an infinite component, compute $Cl_0(x), \dots, Cl_{r-1}(x)$.
 - (a) Make state q_{ip+x}^{jp+y-x} final if y^{i+j} is in an infinite component and $y \in Cl_j(x)$.

- (b) Make state $\check{q}_{ip+x}^{jp+y-x}$ final if $y \in Cl_j(x)$
- (4) If x^p is in a finite component, run the algorithm `FiniteReach` on input x^p and make state \hat{q}_x^{jp+y-x} final if $(y, j) \in Q$.
- (5) If x^p is in an infinite component, compute $Cl_0(x), \dots, Cl_{r-1}(x)$. Declare state \tilde{q}_x^{jp+y-x} final if $y \in Cl_j(x)$.

One can show that \mathcal{A}_{Reach} is the desired automaton. To compute the complexity of building \mathcal{A}_{Reach} , we summarize the computation involved.

- (1) For all x^i in $\mathcal{F}^0 \cup \dots \cup \mathcal{F}^p$, decide whether x^i is in a finite component. This involves computing the set $\mathcal{C}[0], \dots, \mathcal{C}[p-1]$ and takes time $O(p^5)$ by Theorem 6.1.
- (2) For all x^i in $\mathcal{F}^0 \cup \dots \cup \mathcal{F}^p$ such that x^i is in a finite component, run `FiniteReach` on input x^i . This takes time $O(p^5)$ by Corollary 7.3.
- (3) For all $x \in V_{\mathcal{F}}$ such that x^p is in an infinite component, compute the sets $Cl_0(x), \dots, Cl_{r-1}(x)$. This requires time $O(p^6)$ by Theorem 7.1.

Therefore the running time required to construct \mathcal{A}_{Reach} is $O(p^6)$. \square

Finally, we present a solution to the **connectivity problem** on unary automatic graphs of finite degree. This result follows from the construction of \mathcal{A}_{Reach} .

Theorem 7.8. *The connectivity problem for unary automatic graph of finite degree \mathcal{G} is solved in $O(p^6)$, where p is the loop constant of the unary automaton recognizing \mathcal{G} .*

Proof. G is a connected graph if and only if \mathcal{A}_{Reach} accepts the language $\{(1^{ip+x}, 1^{jp+y}) \mid 0 \leq x, y < p \text{ and } ip+x \leq jp+y\}$. Therefore, in order to decide if G is connected, the algorithm constructs the automaton \mathcal{A}_{Reach} and checks if all states in \mathcal{A}_{Reach} are final. This takes time $O(p^6)$. \square

8. CONCLUSION

In this paper we addressed algorithmic problems for graphs of finite degree that have unary automata presentations. We provided polynomial time algorithms that solve connectivity, reachability, infinity testing, and infinite component problems. In our future work we plan to improve these algorithms for other stronger classes of unary automatic graphs. We also point out that there are many other algorithmic problems for finite graphs that can be studied for the class of unary automatic graphs. These, for example, may concern finding spanning trees for automatic graphs, studying the isomorphism problems, and other related issues.

REFERENCES

- [1] A. Blumensath, *Automatic Structures*. Diploma Thesis, RWTH Aachen, 1999.
- [2] A. Blumensath, E. Grädel. *Finite presentations of infinite structures: Automata and interpretations*. Theory of Computing Systems, vol. 37, pp. 642-674, 2004.
- [3] A. Bouajjani, J. Esparza, and O. Maler. *Reachability analysis of pushdown automata: Application to model-checking*. In Proceedings of CONCUR'97, LNCS 1243, pages 135-150, 1997.
- [4] J. R. Büchi, *On a decision method in restricted second-order arithmetic*. Proc. International Congress on Logic, Methodology and Philosophy of Science (E. Nagel, P. Suppes, A. Tarski, Eds.), Stanford University Press, 1-11, 1960.
- [5] D. Caucal, *On infinite graphs having a decidable monadic theory*. In Proc. 27th MFCS (K. Diks, W. Rytter, Eds.), Springer LNCS 2420, 165-176, 2002.
- [6] J. Esparza, D. Hansel, P. Rossmanith, S. Schwoon, *Efficient algorithms for model checking pushdown systems*, Proc. CAV 2000, LNCS 1855, Springer-Verlag, 232-247, 2000.
- [7] P. Hell, J. Nešetřil, *Graphs and Homomorphisms*. Oxford University Press, 2004.
- [8] B. Khoussainov, M. Minnes, *Automatic structures and their complexity*. Submitted to FOCS 2007.
- [9] B. Khoussainov, A. Nerode, *Automatic presentation of structures*. Lecture Notes in Computer Science, 960; 367-392, 1995.
- [10] B. Khoussainov, A. Nies, S. Rubin, F. Stephan, *Automatic structures: richness and limitations*. In Proc. 19th LICS: 44-53, 2004.

- [11] B. Khoussainov, S. Rubin, *Graphs with automatic presentations over a unary alphabet*. Journal of Automata, Languages and Combinatorics 6(4): 467-480,2001.
- [12] B. Khoussainov, S. Rubin, F. Stephan, *Automatic linear orders and trees*. ACM Trans. Comput. Log. 6(4): 675-700, 2005.
- [13] L. Libkin, *Elements of finite model theory*. Springer-Verlag, 2004.
- [14] M. Lohrey, *Automatic structures of bounded degree*. Proc. 10th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR), LNAI 2850: 344 - 358, 2003.
- [15] G. P. Oliver, R. M. Thomas, *Automatic presentations for finitely generated groups*. in Proc. 22nd STACS (V. Diekert, B.Durand, Eds.), Springer, LNCS 3404: 693 - 704, 2005.
- [16] M.O. Rabin, *Decidability of second-order theories and automata on infinite trees*, Trans. Amer. Math. Soc. 141, 1-35, 1969.
- [17] S. Rubin, *Automatic Structures*, PhD Thesis, University of Auckland, 2004.
- [18] W. Thomas, *A short introduction to infinite automata*, In Proceedings of the 5th International Conference Development in Language Theory, Springer, LNCS 2295: 130-144, 2002.