QUERY LEARNING OF DERIVED ω -TREE LANGUAGES IN POLYNOMIAL TIME

DANA ANGLUIN^a, TIMOS ANTONOPOULOS^a, AND DANA FISMAN^b

^a Yale University, New Haven, CT, USA *e-mail address*: {dana.angluin,timos.antonopoulos}@yale.edu

^b Ben-Gurion University, Be'er Sheva, Israel *e-mail address*: dana@cs.bgu.ac.il

ABSTRACT. We present the first polynomial time algorithm to learn nontrivial classes of languages of infinite trees. Specifically, our algorithm uses membership and equivalence queries to learn classes of ω -tree languages derived from weak regular ω -word languages in polynomial time. The method is a general polynomial time reduction of learning a class of derived ω -tree languages to learning the underlying class of ω -word languages, for any class of ω -word languages recognized by a deterministic Büchi acceptor. Our reduction, combined with the polynomial time learning algorithm of Maler and Pnueli [MP95] for the class of weak regular ω -word languages yields the main result. We also show that subset queries that return counterexamples can be implemented in polynomial time using subset queries that return no counterexamples for deterministic or non-deterministic finite word acceptors, and deterministic or non-deterministic Büchi ω -word acceptors.

A previous claim of an algorithm to learn regular ω -trees due to Jayasrirani, Begam and Thomas [JBT08] is unfortunately incorrect, as shown in [Ang16].

1. INTRODUCTION

Query learning is a framework in which a learning algorithm attempts to identify a target concept using specified types of queries to an oracle (or teacher) about the target concept [Ang88]. For example, if the target concept is a regular language L, a membership query asks whether a string x is a member of L, and is answered either "yes" or "no". An equivalence query asks whether a hypothesis language L' (represented, for example, by a deterministic finite acceptor) is equal to L. In the case of an equivalence query, the answer may be "yes", in which case the learning algorithm has succeeded in exact identification of the target concept, or it may be "no", accompanied by a counterexample, that is, a string xin L but not in L' (or vice versa). The counterexample is a witness that L' is not equal to L.

When L' is not equal to L, there is generally a choice (often an infinity) of possible counterexamples, and we require that the learning algorithm works well regardless of which counterexample is chosen by the teacher. To account for this in terms of quantifying the running time of the learning algorithm, we include a parameter that is the maximum length of any counterexample returned by the teacher at any point in the learning process. In this setting, the L^* algorithm of Angluin [Ang87] learns any regular language L using membership



DOI:10.23638/LMCS-15(3:21)2019

and equivalence queries in time polynomial in the size of the smallest deterministic finite acceptor for L and the length of the longest counterexample chosen by the teacher. As shown in [Ang90], there can be no such polynomial time algorithm using just membership queries or just equivalence queries.

The assumption that equivalence queries are available may seem unrealistic. How is a person or a program to judge the equivalence of the target concept to some large, complex, technical specification of a hypothesis? If the hypothesis and the target concept are both deterministic finite acceptors, there is a polynomial time algorithm to test equivalence and return a counterexample in case the answer is negative. Alternatively, if there is a polynomial time algorithm for exact learnability of a class $\mathbb C$ of concepts using membership and equivalence queries, then it may be transformed into a polynomial time algorithm that learns approximations of concepts from $\mathbb C$ using membership queries and randomly drawn labeled examples [Ang87, Ang88]. In this transformation, there is an unknown probability distribution on examples, and an approximation bound $\epsilon > 0$ and a confidence bound $\delta > 0$ are given, and the algorithm draws a corpus of labeled examples of cardinality polynomial in the size of the target concept, $1/\epsilon$ and $\log(1/\delta)$. To answer an equivalence query, the hypothesis is checked against the labeled examples in the corpus. If the hypothesis agrees with the labels of all the examples in the corpus, the equivalence query is answered "yes". and otherwise, any exception supplies a counterexample to return as the answer of the equivalence query. The final hypothesis output by the transformed algorithm will, with probability at least $1-\delta$, have a probability of at most ϵ of disagreeing with the target on examples drawn from the unknown probability distribution.

Since the publication of L^* , there have been a number of substantial improvements and extensions of the algorithm, as well as novel and unanticipated applications in the analysis, verification and synthesis of programs, protocols and hardware, following the work of Peled et al. that identified the applicability of L^* in the area of formal methods [PVY02]. In a recent CACM review article, Vaandrager [Vaa17] explains Model Learning, which takes a black box approach to learning a finite state model of a given hardware or software system using membership queries (implemented by giving the system a sequence of inputs and observing the sequence of outputs) and equivalence queries (implemented using a set of test sequences in which the outputs of the hypothesis are compared with the outputs of the given system.) The learned models may then be analyzed to find discrepancies between a specification and its implementation, or between different implementations. He cites applications in telecommunications [HHNS02, SG14], the automotive industry [FLM⁺13], online conference systems [WNS⁺13], as well as analyzing botnet protocols [CBSS10], smart card readers [CPPdR14], bank card protocols [AdRP13], network protocols [dRP15] and legacy software [MNRS04, SHV16].

Another application of finite state machine learning algorithms is in the assume-guarantee approach to verifying systems by dividing them into modules that can be verified individually. Cobleigh, Giannakopoulou and Păsareănu [CGP03] first proposed using a learning algorithm to learn a correct and sufficient contextual assumption for the component being verified, and there has since been a great deal of research progress in this area [NA06].

If we consider *reactive systems*, that is, systems that maintain an ongoing interaction with their environment (e.g., operating systems, communication protocols, or robotic swarms), the restriction to models specified by finite automata processing finite sequences of inputs is too limiting. Instead, one trajectory of the behavior of a reactive system may be modeled using an infinite word (ω -word), each symbol of which specifies the current state of the

system and the environment at a given time. The system itself may be modeled by an ω -automaton, that is, a finite state automaton that processes ω -words. The desired behavior of such a system may be specified by a linear temporal logic formula, that defines the set of ω -words that constitute "good" behaviors of the system.

Researchers have thus sought query learning algorithms for ω -automata that could be used in the settings of model learning or assume-guarantee verification for reactive systems. However, learning ω -automata seems to be a much more challenging problem than learning automata on finite words, in part because the Myhill-Nerode characterization for regular languages (stating that there is a unique minimum deterministic acceptor that can be constructed using the right congruence classes of the language) does not hold in general for regular ω -languages. The Myhill-Nerode characterization is the basis of the L^* algorithm and its successors.

There is no known polynomial time algorithm using membership and equivalence queries to learn even the whole class \mathbb{DBW} of languages recognized by deterministic Büchi acceptors, which is a strict subclass of the class of all regular ω -languages. Maler and Pnueli [MP95] have given a polynomial time algorithm using membership and equivalence queries to learn the weak regular ω -languages. This class, denoted \mathbb{DwPW} , is the set of languages accepted by deterministic weak parity automata, and is a non-trivial subclass of \mathbb{DBW} . The class \mathbb{DwPW} does have a Myhill-Nerode property, but this alone does not suffice for extending L^* to learn this class, since the observed data might suggest conflicting ways to mark accepting states in an automaton agreeing with the observed data. Maler and Pnueli's algorithm manages to overcome this problem by finding a set of membership queries to ask to resolve the conflict.

In the context of assume-guarantee verification, Farzan et al. [FCC⁺08] proposed a direct application of L^* to learn the full class of regular ω -languages. Their approach is based on the result of Calbrix, Nivat and Podelski [CNP94] showing that a regular ω -language L can be characterized by the regular language $L_{\$}$ of finite strings u\$v representing the set of ultimately periodic words $u(v)^{\omega}$ in L. This establishes that a regular ω -language L is learnable using membership and equivalence queries in time polynomial in the size of the minimal deterministic finite acceptor for $L_{\$}$. However, the size of this representation may be exponentially larger than its ω -automaton representation. More recently, Angluin and Fisman [AF16] have given a learning algorithm using membership and equivalence queries for general regular ω -languages represented by families of deterministic finite acceptors, which improves on the $L_{\$}$ representation, however the running time is not bounded by a polynomial in the representation. Clearly, much more research is needed in the area of query learning of ω -automata.

Despite the difficulties in learning ω -automata, which are used in the analysis of linear temporal logic, in this paper we consider the theoretical question of learning ω -tree automata, which are used in the analysis of branching temporal logic [ES88, KVW00]. As a potential motivation for studying learning of ω -tree languages, we consider a setting in which two players play an infinite game in which the opponent chooses one of two actions (1 and 2) and the player responds with a symbol chosen from a finite alphabet Σ . We can represent the strategy of the player as a binary ω -tree in which each node is the player's state, the two edges leaving the node are the possible choices of the opponent (action 1 or 2), each edge is labeled with the response action (from Σ) of the player, and each leads to a (potentially new) state for the player. In this interpretation, a set of ω -trees represents a property of strategies, and the task of the learner is to learn an initially unknown property of strategies by using membership queries ("Does this strategy have the unknown property?") and equivalence queries ("Is this property the same as the unknown property of strategies?") answered either "yes" or with a counterexample, that is, a strategy that distinguishes the two properties.

Because of the difficulty of the problem, we restrict our attention to ω -tree languages such that all of their paths satisfy a certain temporal logic formula, or equivalently, a property of ω -words. Given an ω -word language L, we use $Trees_d(L)$ to denote the set of all *d*-ary ω -trees t all of whose paths are in L. The ω -tree language $Trees_d(L)$ is often referred to as the *derived* language of L. In this context, it is natural to ask whether learning a derived ω -tree language $Trees_d(L)$ can be reduced to learning the ω -word language L.

We answer this question affirmatively for the case that L can be recognized by a deterministic Büchi word automaton and learned using membership and equivalence queries. Applying this reduction to the result of Maler and Pnueli on polynomially learning languages in $\mathbb{D} \mathbb{W} \mathbb{P} \mathbb{W}$ we obtain a polynomial learning algorithm for derived languages in $Trees_d(\mathbb{D} \mathbb{W} \mathbb{P} \mathbb{W})$ using membership and equivalence queries. Moreover, any progress on polynomially learning an extended subclass \mathbb{C} of $\mathbb{D} \mathbb{B} \mathbb{W}$ using membership and equivalence queries can be automatically lifted to learning $Trees_d(\mathbb{C})$.

The framework of the reduction is depicted in Fig. 1. An algorithm $\mathbf{A}_{\text{Trees}}$ for learning $\text{Trees}_d(L)$ uses a learning algorithm \mathbf{A} for L to complete its task. The membership and equivalence queries (MQ and EQ, henceforth) of algorithm $\mathbf{A}_{\text{Trees}}$ are answered by respective oracles MQ and EQ for $\text{Trees}_d(L)$. Since \mathbf{A} asks membership and equivalence queries about L rather than $\text{Trees}_d(L)$, the learner $\mathbf{A}_{\text{Trees}}$ needs to find a way to answer these queries. If \mathbf{A} asks a membership query about an ω -word, $\mathbf{A}_{\text{Trees}}$ can ask a membership query about an ω -word. Since the tree is accepted by $\text{Trees}_d(L)$ iff the given word is accepted by L it can pass the answer as is to \mathbf{A} . If \mathbf{A} asks an equivalence query, using an acceptor M for an ω -language, then $\mathbf{A}_{\text{Trees}}$ can ask an equivalence query is answered positively then $\mathbf{A}_{\text{Trees}}$ can output the tree accepted by M. If this query is answered positively then $\mathbf{A}_{\text{Trees}}$ can output the tree acceptor M^T and halt. The challenge starts when this query is answered negatively.

When the EQ is answered negatively, a counterexample tree t is given. There are two cases to consider. Either this tree is in $Trees_d(L)$ but is rejected by the hypothesis acceptor M^T , in which case t is referred to as a *positive counterexample*; or this tree is not in $Trees_d(L)$ but is accepted by the hypothesis acceptor M^T , in which case t is referred to as a *negative counterexample*. If t is a positive counterexample, since M^T rejects t there must be a path in t which is rejected by M. It is not too dificult to extract that path. The real challenge is dealing with a negative counterexample. This part is grayed out in the figure. In this case the tree t is accepted by M^T yet it is not in $Trees_d(L)$. Thus, all the paths of the tree are accepted by M, yet at least one path is not accepted by L. Since L is not given, it is not clear how we can extract such a path. Since we know that not all paths of t are contained in L, a use of an unrestricted subset query could help us. Unrestricted subset queries (USQ) are queries on the inclusion of a current hypothesis in the unknown language that are answered by "yes" or "no" with an accompanying counterexample in the case the answer is negative.

Since we don't have access to USQs we investigate whether we can obtain such queries given the queries we have. We show that unrestricted subset queries can be simulated by restricted subset queries. Restricted subset queries (RSQ) on ω -words are subset queries that are *not* accompanied by counterexamples. This essentially means that there is a way to construct a desired counterexample without it being given. To discharge the use of restricted subset queries (as the learner is not provided such queries either) we investigate the relation between subsets of ω -words and ω -trees. Finally, we show that the desired subset queries



FIGURE 1. The reduction framework

on ω -words can be given to the ω -tree learning algorithm by means of ω -tree membership queries. From these we can construct a series of procedures to implement the grayed area.

The subsequent sections contain definitions of ω -words, ω -trees and automata processing them, derived ω -tree languages, the problem of learning classes of ω -word and ω -tree languages, preliminary results, the algorithm for the main reduction, and some discussion. We also include an Appendix with a few examples illustrating some of the procedures involved in our framework.

2. Definitions

2.1. Words and trees. (For more details see Grädel, Thomas and Wilke [GTW02], Perrin and Pin [PP04], and Löding [Löd11].) Let Σ be a fixed finite alphabet of symbols. The set of all finite words over Σ is denoted Σ^* . The empty word is denoted ε , and the length of a finite word x is denoted |x|. Σ^+ is the set of all nonempty finite words over Σ , and for a nonnegative integer k, Σ^k is the set of all finite words over Σ of length equal to k. A finite word language is a subset of Σ^* .

An ω -word over Σ is an infinite sequence $w = \sigma_1 \sigma_2 \sigma_3 \cdots$ where each $\sigma_i \in \Sigma$. The set of all ω -words over Σ is denoted Σ^{ω} . An ω -word language is a subset of Σ^{ω} . The ω -regular expressions are analogous to finite regular expressions, with the added operation S^{ω} , where S is a set of finite words, and the restriction that concatenation combines a set of finite words as the left argument with a set of finite words or ω -words as the right argument. The set S^{ω} is the set of all ω -words $s_1 s_2 \cdots$ such that for each $i, s_i \in S$ and $s_i \neq \varepsilon$. For example, $(a + b)^*(a)^{\omega}$ is the set of all ω -words over $\{a, b\}$ that contain finitely many occurrences of b.

If $S \subseteq \Sigma^*$, *n* is a nonnegative integer and $u \in \Sigma^*$, we define the *length and prefix* restricted version of S by $S[n, u] = S \cap \Sigma^n \cap (u \cdot \Sigma^*)$. This is the set of all words in S of length n that begin with the prefix u. We also define the *length restricted* version of S by $S[n] = S[n, \varepsilon]$, that is, the set of all words in S of length n.

Let d be a positive integer. We consider T_d , the unlabeled complete d-ary ω -tree whose directions are specified by $D = \{1, \ldots, d\}$. The nodes of T_d are the elements of D^* . The root of T_d is the node ε , and the *children* of node v are $v \cdot i$ for $i \in D$. An *infinite path* π in T_d is a sequence x_0, x_1, x_2, \ldots of nodes of T_d such that x_0 is the root and for all nonnegative integers n, x_{n+1} is a child of x_n . An infinite path in T_d corresponds to an ω -word over Dgiving the sequence of directions traversed by the path starting at the root.

A labeled *d*-ary ω -tree (or just ω -tree) is given by a mapping $t: D^+ \to \Sigma$ that assigns a symbol in Σ to each non-root node of T_d . We may think of t as assigning the symbol $t(v \cdot i)$ to the tree edge from node v to its child node $v \cdot i$. The set of all labeled *d*-ary ω -trees is denoted T_d^{Σ} . An ω -tree language is a subset of T_d^{Σ} . If $\pi = x_0, x_1, x_2, \ldots$ is an infinite path of T_d , then we define $t(\pi)$ to be the ω -word $t(x_1), t(x_2), \ldots$ consisting of the sequence of labels of the non-root nodes of π in t. (Recall that t does not label the root node.)

2.2. Automata on words. A finite state word automaton is given by a tuple $M = (Q, q_0, \delta)$, where Q is a finite set of states, $q_0 \in Q$ is the initial state, and $\delta : Q \times \Sigma \to 2^Q$ is the (nondeterministic) transition function. The automaton is *deterministic* if $\delta(q, \sigma)$ contains at most one state for every $(q, \sigma) \in Q \times \Sigma$, and *complete* if $\delta(q, \sigma)$ contains at least one state for every $(q, \sigma) \in Q \times \Sigma$. For a complete deterministic automaton we extend δ to map $Q \times \Sigma^*$ to Q in the usual way.

Let $x = \sigma_1 \sigma_2 \cdots \sigma_k$ be a finite word, where each $\sigma_n \in \Sigma$. A run of M on x is a sequence of k + 1 states r_0, r_1, \ldots, r_k such that $r_0 = q_0$ is the initial state and $r_n \in \delta(r_{n-1}, \sigma_n)$ for integers $1 \le n \le k$. Let $w = \sigma_1 \sigma_2 \cdots$ be an ω -word, where each $\sigma_n \in \Sigma$. A run of M on w is an infinite sequence of states r_0, r_1, r_2, \ldots such that $r_0 = q_0$ is the initial state and $r_n \in \delta(r_{n-1}, \sigma_n)$ for all positive integers n.

A nondeterministic finite acceptor is given by $M = (Q, q_0, \delta, F)$, where (Q, q_0, δ) is a finite state word automaton, and the new component $F \subseteq Q$ is the set of accepting states. M is a deterministic finite acceptor if δ is deterministic. Let M be a nondeterministic finite acceptor and $x \in \Sigma^*$ a finite word of length n. M accepts x iff there is a run r_0, r_1, \ldots, r_n of M on x such that $r_n \in F$. The language recognized by M is the set of all finite words accepted by M, denoted by [M]. The class of all finite word languages recognized by deterministic finite acceptors is denoted by \mathbb{DFW} , and by nondeterministic finite acceptors, \mathbb{NFW} . These representations are equally expressive, that is, $\mathbb{NFW} = \mathbb{DFW}$.

Turning to finite state word automata processing ω -words, a variety of different acceptance criteria have been considered. Such an acceptor is given by a tuple $M = (Q, q_0, \delta, \alpha)$, where (Q, q_0, δ) is a finite state word automaton and α specifies a mapping from 2^Q to $\{0, 1\}$ which gives the criterion of acceptance for an ω -word w.

Given an ω -word w and a run $r = r_0, r_1, \ldots$ of M on w, we consider the set Inf(r) of all states $q \in Q$ such that $r_n = q$ for infinitely many indices n. The acceptor M accepts the ω -word w iff there exists a run r of M on w such that $\alpha(Inf(r)) = 1$. That is, M accepts wiff there exists a run of M on w such that the set of states visited infinitely often in the run satisfies the acceptance criterion α . The language recognized by M is the set of all ω -words accepted by M, denoted [M].

For a Büchi acceptor, the acceptance criterion α is specified by giving a set $F \subseteq Q$ of accepting states and defining $\alpha(S) = 1$ iff $S \cap F \neq \emptyset$. In words, a Büchi acceptor M accepts w if and only if there exists a run r of M on w such that at least one accepting state is

visited infinitely often in r. For a co-Büchi acceptor, the acceptance criterion α is specified by giving a set $F \subseteq Q$ of rejecting states and defining $\alpha(S) = 1$ iff $S \cap F = \emptyset$. For a parity acceptor, α is specified by giving a function c mapping Q to an interval of integers [i, j], (called *colors* or *priorities*) and defining $\alpha(S) = 1$ iff the minimum integer in c(S) is even.

A parity automaton is said to be weak if no two strongly connected states have distinct colors, i.e., if looking at the partition of its states to maximal strongly connected components (MSCCs) all states of an MSCC have the same color. Clearly every weak parity automaton can be colored with only two colors, one even and one odd, in which case the colors are often referred to as accepting or rejecting. It follows that a weak parity automaton can be regarded as either a Büchi or a coBüchi automaton. If in addition no rejecting MSCC is reachable from an accepting MSCC, the acceptor is said to be weak Büchi. Likewise, a weak parity acceptor where no accepting MSCC is reachable from a rejecting MSCC, is said to be weak coBüchi acceptor.

The classes of languages of ω -words recognized by these kinds of acceptors will be denoted by three/four-letter acronyms, with \mathbb{N} or \mathbb{D} (for nondeterministic or deterministic), \mathbb{B} , \mathbb{C} , \mathbb{P} , \mathbb{wB} , \mathbb{wC} or \mathbb{wP} (for Büchi, co-Büchi, parity or their respective weak variants) and then \mathbb{W} (for ω -words). Thus $\mathbb{D}\mathbb{wB}\mathbb{W}$ is the class of ω -word languages recognized by deterministic weak Büchi word acceptors.

Concerning the expressive power of various types of acceptors, previous research has established the following results. The weak variants are strictly less expressive than the non-weak variants. Deterministic parity automata are more expressive than deterministic Büchi and coBüchi automata and the same is true for their weak variants. These results are summarized in Fig. 2. In addition, $\mathbb{NBW} = \mathbb{DPW} = \mathbb{NPW}$ and $\mathbb{DwPW} = \mathbb{DCW} \cap \mathbb{DBW}$. The class of regular ω -languages is the class \mathbb{DPW} , and the class of weak regular ω -languages is the class \mathbb{DwPW} .



FIGURE 2. Expressiveness hierarchy of ω -acceptors [Wag75, MP90].

2.3. Automata on trees. Acceptors on *d*-ary ω -trees are equipped with analogous accepting conditions. Such an acceptor is given by a tuple $M = (Q, q_0, \delta, \alpha)$, where Q is a finite set of states, $q_0 \in Q$ is the initial state, the transition function δ is a map from Q and *d*-tuples of symbols to sets of *d*-tuples of states, that is, $\delta : Q \times \Sigma^d \to 2^{Q^d}$, and the acceptance criterion α specifies a function from 2^Q to $\{0, 1\}$. We may think of the acceptor as running top down from the root of the tree, at each node nondeterministically choosing a permissible *d*-tuple of states for the *d* children of the node depending on the state assigned to the node and the *d*-tuple of symbols on its outgoing edges. In other words, for each node, with a state qassigned to it, and *d* outgoing edges with symbols $\sigma_1, \ldots, \sigma_d$, the acceptor will assign states q_1, \ldots, q_d to the children of the nodes, only if $(q_1, \ldots, q_d) \in \delta(q, (\sigma_1, \ldots, \sigma_d))$.

We define a run of M on the ω -tree t as a mapping r from the nodes of T_d to Q such that $r(\varepsilon) = q_0$ and for every node x, we have $(r(x \cdot 1), \ldots, r(x \cdot d)) \in \delta(r(x), (t(x \cdot 1), \ldots, t(x \cdot d)))$. That is, the root is assigned the initial state and for every node, the ordered d-tuple of states assigned to its children is permitted by the transition function. The acceptor M accepts the ω -tree t iff there exists a run r of M on t such that for every infinite path π , we have $\alpha(Inf(r(\pi))) = 1$. That is, there must be at least one run in which, for every infinite path, the set of states that occur infinitely often on the path satisfies the acceptance criterion α . The ω -tree language recognized by M is the set of all ω -trees accepted by M, denoted [M]. The specification of the acceptance criterion α is as for ω -word acceptors, yielding Büchi, co-Büchi and parity ω -tree acceptors. If the transition function specifies at most one permissible *d*-tuple of states for every element of $Q \times \Sigma^d$, then the acceptor is deterministic. The corresponding classes of ω -tree languages are also denoted by three-letter acronyms, where the last letter is \mathbb{T} for ω -trees. For ω -trees, the class of all regular ω -tree languages is \mathbb{NPT} and \mathbb{NBT} is a proper subclass of \mathbb{NPT} . For any automaton or acceptor M, we denote the number of its states by |M|.

3. Derived ω -tree languages

Given an ω -tree t we define the ω -word language paths(t) consisting of the ω -words labeling its infinite paths. That is, we define

$$paths(t) = \{t(\pi) \mid \pi \text{ is an infinite path in } T_d\}.$$

If L is an ω -word language and d is a positive integer, we define a corresponding language of d-ary ω -trees derived from L as follows:

$$Trees_d(L) = \{ t \in T_d^{\Sigma} \mid paths(t) \subseteq L \}.$$

That is, $Trees_d(L)$ consists of all *d*-ary ω -trees such that every infinite path in the tree is labeled by an element of *L*. If \mathbb{C} is any class of ω -word languages, $Trees_d(\mathbb{C})$ denotes the class of all ω -tree languages $Trees_d(L)$ such that $L \in \mathbb{C}$.

3.1. **Derived tree languages.** Not every regular *d*-ary ω -tree language can be derived in this way from an ω -word language. As an example, consider the language L_a of all binary ω -trees t over $\Sigma = \{a, b\}$ such that there is at least one node labeled with a. An NBT acceptor can recognize L_a by guessing and checking a path that leads to an a. However, if $L_a = \text{Trees}_2(L)$ for some ω -word language L, then because there are ω -trees in L_a that have infinite paths labeled exclusively with b, we must have $b^{\omega} \in L$, so the binary ω -tree labeled exclusively with b would also be in $\text{Trees}_2(L)$, a contradiction.

Given an ω -word acceptor $M = (Q, q_0, \delta, \alpha)$, we may construct a related ω -tree acceptor $M^{T,d} = (Q, q_0, \delta^{T,d}, \alpha)$ as follows. For all $q \in Q$ and all $(\sigma_1, \ldots, \sigma_d) \in \Sigma^d$, define

$$\delta^{T,d}(q,(\sigma_1,\ldots,\sigma_d)) = \{(q_1,\ldots,q_d) \mid \forall i \in D, q_i \in \delta(q,\sigma_i)\}$$

That is, the acceptor $M^{T,d}$ may continue the computation at a child of a node with any state permitted by M, independently chosen. It is tempting to think that $\llbracket M^{T,d} \rrbracket = Trees_d(\llbracket M \rrbracket)$, but this may not be true when M is not deterministic.

Lemma 3.1. Given an ω -word acceptor M, we have that $\llbracket M^{T,d} \rrbracket \subseteq \operatorname{Trees}_d(\llbracket M \rrbracket)$ with equality if M is deterministic.

Proof. Consider the ω -word acceptor $M = (Q, q_0, \delta, \alpha)$. If $t \in \llbracket M^{T,d} \rrbracket$ then there is a run r of $M^{T,d}$ on t satisfying the acceptance criterion α on every infinite path. Thus, $t(\pi) \in \llbracket M \rrbracket$ for every infinite path π and $t \in Trees_d(\llbracket M \rrbracket)$.

Suppose $t \in Trees_d(\llbracket M \rrbracket)$ and M is deterministic. Then $M^{T,d}$ is also deterministic, and there is a unique run r of $M^{T,d}$ on t. For every infinite path π , $r(\pi)$ is also the unique run of M on the ω -word $t(\pi)$, which satisfies α because $t \in Trees_d(\llbracket M \rrbracket)$. Thus $t \in \llbracket M^{T,d} \rrbracket$. \Box Boker et al. [BKKS13] give the following example to show that the containment asserted in Lemma 3.1 may be proper if M is not deterministic. The ω -language L specified by $(a + b)^* b^{\omega}$ can be recognized by the nondeterministic Büchi acceptor M with two states, q_0 and q_1 , transition function $\delta(q_0, a) = \{q_0\}$, $\delta(q_0, b) = \{q_0, q_1\}$, $\delta(q_1, b) = \{q_1\}$, and accepting state set $\{q_1\}$. Let d = 2, specifying binary trees with directions $\{1, 2\}$. Then $M^{T,2}$ is a nondeterministic ω -tree acceptor, but the following example shows $[\![M^{T,2}]\!] \subseteq Trees_2(L)$. Consider the binary ω -tree t that labels every node in 1*2 with a and every other non-root node with b. Clearly $t \in Trees_2(L)$ because every infinite path in t has at most one a, but no run of $M^{T,2}$ can satisfy the acceptance criterion on the path 1 $^{\omega}$. Suppose r were an accepting run of $M^{T,2}$ on t. Then for some $n \ge 0$, $r(1^n)$ would have to be equal to q_1 . But then such a mapping r would not be a valid run because $1^n 2$ is labeled by a and $\delta^{T,2}(q_1, (b, a)) = \emptyset$ because $\delta(q_1, a) = \emptyset$.

3.2. Good for trees. This phenomenon motivates the following definition. An ω -word acceptor M is good for trees iff for any positive integer d, $\llbracket M^{T,d} \rrbracket = Trees_d(\llbracket M \rrbracket)$. Nondeterministic ω -word acceptors that are good for trees are equivalent in expressive power to deterministic ω -word acceptors, as stated by the following result of Boker et al.

Theorem 3.2 [BKKS13]. Let L be a regular ω -word language and $d \ge 2$. If $\operatorname{Trees}_d(L)$ is recognized by a nondeterministic ω -tree acceptor with acceptance criterion α , then L can be recognized by a deterministic ω -word acceptor with acceptance criterion α .

This theorem generalizes prior results of Kupferman, Safra and Vardi for Büchi acceptors [KSV06] and Niwiński and Walukiewicz for parity acceptors [NW98]. One consequence of Theorem 3.2 is that when $d \ge 2$, nondeterministic ω -word acceptors that are good for trees are not more expressive than the corresponding deterministic ω -word acceptors. Also, for $d \ge 2$, nondeterminism does not increase expressive power over determinism when recognizing ω -tree languages of the form $Trees_d(L)$. To see this, if N is a nondeterministic ω -tree acceptor with acceptance criterion α recognizing $Trees_d(L)$ then there is a deterministic ω -word acceptor M with acceptance criterion α such that $[\![M]\!] = L$, and $M^{T,d}$ is a deterministic ω -tree acceptor with acceptance criterion α that also recognizes $Trees_d(L)$.

However, it is possible that nondeterminism permits acceptors with smaller numbers of states. Kuperberg and Skrzypczak [KS15] have shown that for an NBT acceptor Mrecognizing the ω -tree language $Trees_d(L)$, there is a DBW acceptor with at most $|M|^2$ states recognizing L, so nondeterminism gives at most a quadratic savings for Büchi tree acceptors that are good for trees. However, they have also shown that the blowup in the case of nondeterministic co-Büchi tree acceptors (and all higher parity conditions) is necessarily exponential in the worst case.

4. LEARNING TREE LANGUAGES

We address the problem of learning derived ω -tree languages by giving a polynomial time reduction of the problem of learning $Trees_d(\mathbb{C})$ to the problem of learning \mathbb{C} . The paradigm of learning we consider is exact learning with membership queries and equivalence queries. Maler and Pnueli [MP95] have given a polynomial time algorithm to learn the class of weak regular ω -languages using membership and equivalence queries. Their algorithm and the reduction we give in Theorem 7.12 prove the following theorem. **Theorem 4.1.** For every positive integer d, there is a polynomial time algorithm to learn $\operatorname{Trees}_d(\mathbb{D} \rtimes \mathbb{P} \mathbb{W})$ using membership and equivalence queries.

4.1. **Representing examples.** For a learning algorithm, the examples tested by membership queries and the counterexamples returned by equivalence queries need to be finitely represented. For learning regular ω -word languages, it suffices to consider *ultimately periodic* ω -words, that is, words of the form $u(v)^{\omega}$ for finite words $u \in \Sigma^*$ and $v \in \Sigma^+$. If two regular ω -word languages agree on all the ultimately periodic ω -words, then they are equal. The pair (u, v) of finite words represents the ultimately periodic word $u(v)^{\omega}$.

The corresponding class of examples in the case of regular ω -tree languages is the class of regular ω -trees. These are ω -trees that have a finite number of nonisomorphic complete infinite subtrees. We represent a regular ω -tree t by a regular ω -tree automaton $A_t = (Q, q_0, \delta, \tau)$, where (Q, q_0, δ) is a complete deterministic finite state word automaton over the input alphabet $D = \{1, \ldots, d\}$ and τ is an output function that labels each transition with an element of Σ . That is, $\tau : Q \times D \to \Sigma$. The regular ω -tree t represented by such an automaton A_t is defined as follows. For $x \in D^+$, let $i \in D$ be the last symbol of x and let x' be the rest of x, so that $x = x' \cdot i$. Then define $t(x) = \tau(\delta(q_0, x'), i)$, that is, t(x) is the label assigned by τ to the last transition in the unique run of A_t on x.

Rabin [Rab72] proved that if two regular ω -tree languages agree on all the regular ω -trees then they are equal. Thus, ultimately periodic ω -words and regular ω -trees are proper subsets of examples that are nonetheless sufficient to determine the behavior of regular ω -word and ω -tree acceptors on all ω -words and ω -trees, respectively.

4.2. Types of queries for learning. We consider the situation in which a learning algorithm **A** is attempting to learn an initially unknown target language L of ω -words from a known class $\mathbb{C} \subseteq \mathbb{DBW}$. The information that **A** gets about L is in the form of answers to queries of specific types [Ang88]. The learning algorithm will use membership and equivalence queries, whereas, restricted and unrestricted subset queries will in addition be considered in the proof.

In a membership query about L, abbreviated MQ, the algorithm **A** specifies an example as a pair of finite words (u, v) and receives the answer "yes" if $u(v)^{\omega} \in L$ and "no" otherwise. In an equivalence query about L, abbreviated EQ, the algorithm **A** specifies a hypothesis language $[\![M]\!]$ as a DBW acceptor M, and receives either the answer "yes" if $L = [\![M]\!]$, or "no" and a counterexample, that is, a pair of finite words (u, v) such that $u(v)^{\omega} \in (L \oplus [\![M]\!])$, where $B \oplus C$ denotes the symmetric difference of sets B and C.

In a restricted subset query about L, abbreviated RSQ, the algorithm **A** specifies a hypothesis language $[\![M]\!]$ as a DBW acceptor M, and receives the answer "yes" if $[\![M]\!] \subseteq L$ and "no" otherwise. An unrestricted subset query about L, abbreviated USQ, is like a restricted subset query, except that in addition to the answer of "no", a counterexample (u, v) is provided such that $u(v)^{\omega} \in ([\![M]\!] \setminus L)$.

A learning algorithm **A** using specific types of queries *exactly learns* a class \mathbb{C} of ω -word languages iff for every $L \in \mathbb{C}$, the algorithm makes a finite number of queries of the specified types about L and eventually halts and outputs a DBW acceptor M such that [M] = L. The algorithm runs in polynomial time iff there is a fixed polynomial p such that for every $L \in \mathbb{C}$, at every point the number of steps used by **A** is bounded by p(n, m), where n is

the size of the smallest DBW acceptor recognizing L, and m is the maximum length of any counterexample **A** has received up to that point.

The case of a learning algorithm for ω -tree languages is analogous, except that the examples and counterexamples are given by regular ω -tree automata, and the hypotheses provided to equivalence or subset queries are represented by DBT acceptors. We also consider cases in which the inputs to equivalence or subset queries may be NBW or NBT acceptors.

5. FRAMEWORK OF A REDUCTION

Suppose **A** is a learning algorithm that uses membership and equivalence queries and exactly learns a class $\mathbb{C} \subseteq \mathbb{DBW}$. We shall describe an algorithm \mathbf{A}_{Trees} that uses membership and equivalence queries and exactly learns the derived class $Trees_d(\mathbb{C})$ of ω -tree languages. Note that $Trees_d(\mathbb{C}) \subseteq \mathbb{DBT}$.

The algorithm \mathbf{A}_{Trees} with target concept $Trees_d(L)$ simulates algorithm \mathbf{A} with target concept L. In order to do so, \mathbf{A}_{Trees} must correctly answer membership and equivalence queries from \mathbf{A} about L by making one or more membership and/or equivalence queries of its own about $Trees_d(L)$. Before describing the algorithm \mathbf{A}_{Trees} we establish some basic results about regular ω -trees.

5.1. Testing acceptance of a regular ω -tree. We describe a polynomial time algorithm $Accepted?(A_t, M)$ that takes as input a regular ω -tree t represented by a regular ω -tree automaton $A_t = (Q_1, q_{0,1}, \delta_1, \tau_1)$ and a DBW acceptor $M = (Q_2, q_{0,2}, \delta_2, F_2)$ and determines whether or not $M^{T,d}$ accepts t. If not, it also outputs a pair (u, v) of finite words such that $u(v)^{\omega} \in (paths(t) \setminus [M]).$

Algorithm 1 : $Accepted?(A_t, M)$

```
Require: A_t = (Q_1, q_{0,1}, \delta_1, \tau_1) representing t;
  M = (Q_2, q_{0,2}, \delta_2, F_2), a complete DBW acceptor
Ensure: Return "yes" if M^{T,d} accepts t
  else return "no" and (u, v) with u(v)^{\omega} \in (paths(t) \setminus [M]).
  let Q = Q_1 \times Q_2
  let q_0 = (q_{0,1}, q_{0,2})
  for all (q_1, q_2) \in Q and i \in D do
      let \delta((q_1, q_2), i) = (\delta_1(q_1, i), \delta_2(q_2, \tau_1(q_1, i)))
  let F = \{(q_1, q_2) \mid q_2 \in F_2\}
  let M' = (Q, q_0, \delta, F)
  if \llbracket M' \rrbracket = D^{\omega} then
      return "yes"
  else
      find x(y)^{\omega} \in (D^{\omega} \setminus \llbracket M' \rrbracket)
      let u(v)^{\omega} = t(x(y)^{\omega})
       return "no" and (u, v)
```

We may assume M is complete by adding (if necessary) a new non-accepting sink state and directing all undefined transitions to the new state. We construct a DBW acceptor M' over the alphabet $D = \{1, \ldots, d\}$ by combining A_t and M as follows. The states are $Q = Q_1 \times Q_2$, the initial state is $q_0 = (q_{0,1}, q_{0,2})$, the set of accepting states is $F = \{(q_1, q_2) \mid q_2 \in F_2\}$, and the transition function δ is defined by $\delta((q_1, q_2), i) = (\delta_1(q_1, i), \delta_2(q_2, \tau_1(q_1, i)))$ for all $(q_1, q_2) \in Q$ and $i \in D$. For each transition, the output of the regular ω -tree automaton A_t is the input of the DBW acceptor M. An infinite path π in t corresponds to an ω -word $z \in D^{\omega}$, giving the sequence of directions from the root. The unique run of M' on z traverses a sequence of states; if we project out the first component, we get the run of A_t on z, and if we project out the second component, we get the run of M on $t(\pi)$. Then $M^{T,d}$ accepts t iff M accepts $t(\pi)$ for every infinite path π , which is true iff $[\![M']\!] = D^{\omega}$. This in turn is true iff every nonempty accessible recurrent set of states in M' contains at least one element of F.

A set S of states is *recurrent* iff for all $q, q' \in S$, there is a nonempty finite word v such that $\delta(q, v) = q'$ and for every prefix u of v we have $\delta(q, u) \in S$. A set S of states is *accessible* iff for every $q \in S$ there exists a finite word u such that $\delta(q_0, u) = q$.

The algorithm to test whether $M^{T,d}$ accepts t first removes from the transition graph of M' all states that are not accessible. It then removes all states in F and tests whether there is any cycle in the remaining graph. If not, then $M^{T,d}$ accepts t. Otherwise, there is a state q in Q and finite words $x \in D^*$ and $y \in D^+$ such that $\delta(q_0, x) = q$ and $\delta(q, y) = q$ and none of the states traversed from q to q along the path y are in F. Thus, $x(y)^{\omega}$ is an ultimately periodic path π that does not visit F infinitely often, and letting $u(v)^{\omega}$ be $t(x(y)^{\omega})$, we have $u(v)^{\omega} \in (paths(t) \setminus [M])$, so the pair (u, v) is returned in this case. The required graph operations are standard and can be accomplished in time polynomial in |M| and $|A_t|$.

5.2. Representing a language as paths of a tree. When the algorithm \mathbf{A}_{Trees} makes a membership query about $Trees_d(L)$ with a regular ω -tree t, the answer is "yes" if $paths(t) \subseteq L$ and "no" otherwise. Thus, this query has the effect of a restricted subset query about L with paths(t). However, this does not give us restricted subset queries for arbitrary \mathbb{DBW} languages. Next, we examine the close relationship between languages of the form paths(t) and safety languages.

An ω -word language L is a safety language iff L is a regular ω -word language and for every ω -word w not in L, there exists a finite prefix x of w such that no ω -word with prefix x is in L. A language is safety iff it is in the class $\mathbb{D}\mathbb{W}\mathbb{C}\mathbb{W}$. An alternative characterization is that there is an NBW acceptor $M = (Q, q_0, \delta, Q)$, all of whose states are accepting, such that $[\![M]\!] = L$. In this case, the acceptor is typically not complete (otherwise it recognizes Σ^{ω}). An example of a language in $\mathbb{D}\mathbb{W}\mathbb{P}\mathbb{W}$ that is not a safety language is $a^*b^*(a)^{\omega}$. Although b^{ω} is not in the language, every finite prefix b^k is a prefix of some ω -word in the language.

Lemma 5.1. If A_t is a regular ω -tree automaton representing an ω -tree t, then paths(t) is a safety language recognizable by an NBW acceptor M with $|M| = |A_t|$.

Proof. If $A_t = (Q, q_0, \delta, \tau)$, then we define $M = (Q, q_0, \delta', Q)$ where

$$\delta'(q,\sigma) = \{ r \in Q \mid (\exists i \in D) (\delta(q,i) = r \land \tau(q,i) = \sigma) \}$$

for all $q \in Q$ and $\sigma \in \Sigma$. That is, the *M* transition on *q* and σ is defined to be all states reachable from *q* by a transition in A_t labeled with σ . Note that all states of *M* are accepting.

If $w \in paths(t)$, then there is a run r_0, r_1, \ldots of A_t whose transitions are labeled by w, and this is a run of M on w, so $w \in \llbracket M \rrbracket$. Conversely, if $w \in \llbracket M \rrbracket$, then there is some run r_0, r_1, \ldots of M on w, and this is a run of A_t whose transitions are labeled with w, and therefore $w \in paths(t)$.

For the converse, representing a safety language as the paths of a regular ω -tree, we require a lower bound on d, the arity of the tree. If $M = (Q, q_0, \delta, F)$ is an NBW acceptor and $q \in Q$, we define the set of transitions out of q to be transitions $(q) = \{(\sigma, r) \mid \sigma \in Q, \sigma\}$

 $\Sigma \wedge r \in \delta(q, \sigma)$. We define the *out-degree* of M to be the maximum over $q \in Q$ of the cardinality of transitions(q).

Lemma 5.2. Let L be a safety language recognized by NBW acceptor $M = (Q, q_0, \delta, Q)$. Suppose the out-degree of M is at most d. Then there is a d-ary regular ω -tree t such that paths(t) = L, and t is representable by A_t with $|A_t| = |M|$.

Proof. We may assume that every state of M is accessible and has at least one transition defined. We define $A_t = (Q, q_0, \delta_t, \tau)$ over the alphabet $D = \{1, \ldots, d\}$ as follows. For $q \in Q$, choose a surjective mapping f_q from D to transitions(q). Then for $q \in Q$ and $i \in D$, let $(\sigma, r) = f_q(i)$ and define $\delta_t(q, i) = r$ and $\tau(q, i) = \sigma$.

If $w \in L$, then there is a run r_0, r_1, \ldots of M on w, and there is an infinite path in A_t traversing the same states in which the labels are precisely w, so $w \in paths(t)$. Conversely, if $w \in paths(t)$, then there is an infinite path π such that $t(\pi) = w$, and the sequence of states of A_t traversed by w yields a run of M on w, so $w \in L$.

The NBW acceptor in the proof of Lemma 5.1 can be determinized via the subset construction to give a DBW acceptor of size at most $2^{|A_t|}$ recognizing the same language. In the worst case this exponential blow up in converting a regular ω -tree automaton to a DBW acceptor is necessary, as shown by the following lemma.

Lemma 5.3. There exists a family of regular ω -trees t_1, t_2, \ldots such that t_n can be represented by a regular ω -tree automaton of size n + 2, but the smallest DBW acceptor recognizing paths (t_n) has size at least 2^n .

Proof. Let $\Sigma = \{a, b, c\}$ and let L_n be $(a + b + (a(a + b)^n c))^{\omega}$. This is a safety language: $w \in L_n$ iff every occurrence of c in w is preceded by a word of the form $a(a + b)^n$. There is a NBW acceptor M_n of n + 2 states recognizing L_n . The states are nonnegative integers in [0, n + 1], with 0 the initial state, $\delta(0, a) = \{0, 1\}$, $\delta(0, b) = 0$, $\delta(i, a) = \delta(i, b) = i + 1$ for $1 \le i \le n$, and $\delta(n + 1, c) = 0$.

By Lemma 5.2, there is a ternary regular ω -tree t_n such that $paths(t_n) = L_n$ and t_n is represented by a regular ω -tree automaton with n + 2 states. However, any DBW acceptor recognizing L_n must have enough states to distinguish all 2^n strings in $(a + b)^n$ in order to check the safety condition.

If t is a d-ary regular ω -tree represented by the regular ω -tree automaton A_t , then $acceptor(A_t)$ denotes the NBW acceptor M recognizing paths(t) constructed from A_t in the proof of Lemma 5.1. Note that the out-degree of $acceptor(A_t)$ is at most d.

If M is an NBW acceptor such that $\llbracket M \rrbracket$ is a safety language and the out-degree of M is at most d, then $tree_d(M)$ denotes the regular ω -tree automaton A_t constructed from M in the proof of Lemma 5.2. We also use the notation $tree_d(L)$ if L is a safety language and the implied acceptor for L is clear.

For example, given finite words $u \in \Sigma^*$ and $v \in \Sigma^+$, the singleton set containing $u(v)^{\omega}$ is a safety language recognized by a DBW of out-degree 1 and size linear in |u| + |v|. Then $tree_d(u(v)^{\omega})$ represents the *d*-ary tree all of whose infinite paths are labeled with $u(v)^{\omega}$.

6. The algorithm \mathbf{A}_{Trees}

We now describe the algorithm \mathbf{A}_{Trees} , which learns $Trees_d(L)$ by simulating the algorithm \mathbf{A} and answering the membership and equivalence queries of \mathbf{A} about L. It is summarized in Algorithm 2, and some of the cases are illustrated in an example presented in Appendix A.

Algorithm $2 : A_{Trees}$

Require: Learning algorithm A for \mathbb{C} ;
MQ and EQ access to $Trees_d(L)$ for $L \in \mathbb{C}$
Ensure: Acceptor $M^{T,d}$ such that $\llbracket M^{T,d} \rrbracket = Trees_d(L)$
while A has not halted do
if next step of A is not a query then
simulate next step of \mathbf{A}
else if A asks $MQ(u, v)$ about L then
answer A with MQ($tree_d(u(v)^{\omega})$) about $Trees_d(L)$
else if A asks $EQ(M)$ about L then
ask $EQ(M^{T,d})$ about $Trees_d(L)$
if $EQ(M^{T,d})$ answer is "yes" then
return $M^{T,d}$ and halt
else {EQ $(M^{T,d})$ answer is counterexample tree t given by A_t }
if $Accepted?(A_t, M)$ returns "no" with value (u, v) then
answer A with (u, v)
else { $Accepted?(A_t, M)$ returns "yes"}
let $M' = acceptor(A_t)$
for all accepting states q of M' do
simulate in parallel $Findctrex(M',q)$
terminate all computations and answer A with the first (u, v) returned
$[\mathbf{A} \text{ halts with output } M]$
return $M^{T,d}$ and halt

If **A** asks a membership query with (u, v) then \mathbf{A}_{Trees} constructs the regular ω -tree automaton $tree_d(u(v)^{\omega})$ representing the *d*-ary regular ω -tree all of whose infinite paths are labeled $u(v)^{\omega}$, and makes a membership query with $tree_d(u(v)^{\omega})$. Because $u(v)^{\omega} \in L$ iff the tree represented by $tree_d(u(v)^{\omega})$ is in $Trees_d(L)$, the answer to the query about $tree_d(u(v)^{\omega})$ is simply given to **A** as the answer to its membership query about (u, v).

For an equivalence query from **A** specified by a DBW acceptor M, the algorithm \mathbf{A}_{Trees} constructs the corresponding DBT acceptor $M^{T,d}$, which recognizes $Trees_d(\llbracket M \rrbracket)$, and makes an equivalence query with $M^{T,d}$. If the answer is "yes", the algorithm \mathbf{A}_{Trees} has succeeded in learning the target ω -tree language $Trees_d(L)$ and outputs $M^{T,d}$ and halts. Otherwise, the counterexample returned is a regular ω -tree t in $\llbracket M^{T,d} \rrbracket \oplus Trees_d(L)$, represented by a regular ω -tree automaton A_t . A call to the algorithm $Accepted?(A_t, M)$ determines whether $M^{T,d}$ accepts t. If $M^{T,d}$ rejects t, then $t \in Trees_d(L)$ and t is a positive counterexample. If $M^{T,d}$ accepts t, then $t \notin Trees_d(L)$ and t is a negative counterexample. We next consider these two cases.

If t is a positive counterexample then we know that $t \in Trees_d(L)$ and therefore $paths(t) \subseteq L$. Because $t \notin \llbracket M^{T,d} \rrbracket$, the acceptor M must reject at least one infinite path in t. In this case, the algorithm $Accepted?(A_t, M)$ returns a pair of finite words (u, v) such that $u(v)^{\omega} \in (paths(t) \setminus \llbracket M \rrbracket)$, and therefore $u(v)^{\omega} \in (L \setminus \llbracket M \rrbracket)$. The algorithm \mathbf{A}_{Trees} returns the positive counterexample (u, v) to \mathbf{A} in response to its equivalence query with M.

If t is a negative counterexample, that is, $t \in (\llbracket M^{T,d} \rrbracket \setminus \operatorname{Trees}_d(L))$, then we know that $paths(t) \subseteq \llbracket M \rrbracket$, but at least one element of paths(t) is not in L, so $(\llbracket M \rrbracket \setminus L) \neq \emptyset$. Ideally, we would like to extract an ultimately periodic ω -word $u(v)^{\omega} \in (paths(t) \setminus L)$ and provide (u, v) to **A** as a negative counterexample in response to its equivalence query with M.

If we could make an *unrestricted subset query* with paths(t) about L, then the counterexample returned would be precisely what we need.

As noted previously, if t is any regular ω -tree then we can simulate a restricted subset query with paths(t) about L by making a membership query with t about $Trees_d(L)$, because $paths(t) \subseteq L$ iff $t \in Trees_d(L)$. In order to make use of this, we next show how to use restricted subset queries about L to implement an unrestricted subset query about L.

6.1. Restricted subset queries. To establish basic techniques, we show how to reduce unrestricted subset queries to restricted subset queries for nondeterministic or deterministic finite acceptors over finite words. Suppose $L \subseteq \Sigma^*$ and we may ask restricted subset queries about L. In such a query, the input is a nondeterministic (resp., deterministic) finite acceptor M, and the answer is "yes" if [M] is a subset of L, and "no" otherwise. If the answer is "no", we show how to find a shortest counterexample $u \in ([M] \setminus L)$ in time polynomial in |M| and |u|.

Theorem 6.1. There is an algorithm \mathbf{R}^* which takes as input an NFW (resp., DFW) M, and has restricted subset query access to a language L with NFW (resp., DFW) acceptors as inputs, that correctly answers the unrestricted subset query with M about L. Additionally, if L is recognized by a DFW T_L , then $\mathbf{R}^*(M)$ runs in time bounded by a polynomial in |M|and $|T_L|^{1}$

The idea of the proof is to first establish the minimal length ℓ of a counterexample, and then try to extend the prefix ϵ letter by letter until obtaining a full length minimal counterexample. Note that trying to establish a prefix of a counterexample letter by letter, without obtaining a bound first, may not terminate. For instance, if $L = \Sigma^* \setminus a^*b$, one can establish the sequence of prefixes $\epsilon, a, aa, aaa, \ldots$ and never reach a counterexample.

To prove Theorem 6.1 we first construct an acceptor $M_{\ell,v}$ for $[\![M]\!][\ell,v]$, the length and prefix restricted version of $[\![M]\!]$, given M, ℓ and v as inputs.

Lemma 6.2. There is a polynomial time algorithm to construct an acceptor $M_{\ell,v}$ for $[M][\ell, v]$ given a NFW acceptor M, a nonnegative integer ℓ and a finite word v, such that

- (1) $M_{\ell,v}$ has at most one accepting state, which has no out-transitions,
- (2) the out-degree of $M_{\ell,v}$ is at most the out-degree of M,
- (3) $M_{\ell,v}$ is deterministic if M is deterministic.

Proof. If $\ell < |v|$, then $\llbracket M \rrbracket [\ell, v] = \emptyset$, and the output $M_{\ell,v}$ is a one-state acceptor with no accepting states. Otherwise, assume $v = \sigma_1 \sigma_2 \cdots \sigma_k$ and construct M' to be the deterministic finite acceptor for $v \cdot \Sigma^{\ell - |v|}$ with states $0, 1, \ldots, \ell$ where 0 is the initial state, ℓ is the final state, and the transitions are $\delta(i, \sigma_{i+1}) = i + 1$ for $0 \le i < k$ and $\delta(i, \sigma) = i + 1$ for $k \le i < \ell$ and $\sigma \in \Sigma$.

Then $M_{\ell,v}$ is obtained by a standard product construction of M and M' for the intersection $[\![M]\!] \cap [\![M']\!]$, with the observation that no accepting state in the product has any out-transitions defined, so they may all be identified. It is straightforward to verify the required properties of $M_{\ell,v}$.

Proof of Theorem 6.1. For input M, define $M_{[\ell,v]}$ to be the finite acceptor constructed by the algorithm of Lemma 6.2 to recognize the length and prefix restricted language $[\![M]\!][\ell,v]$.

For $\ell = 0, 1, 2, \ldots$, ask a restricted subset query with $M_{[\ell,\varepsilon]}$, until the first query answered "no". At this point, ℓ is the shortest length of a counterexample in $(\llbracket M \rrbracket \setminus L)$. Then a counterexample u of length ℓ is constructed symbol by symbol.

¹The cardinality of Σ is treated as a constant.

Assume we have found a prefix u' of a counterexample of length ℓ in $(\llbracket M \rrbracket \setminus L)$, with $|u'| < \ell$. For each symbol $\sigma \in \Sigma$ we ask a restricted subset query with $M_{[\ell, u'\sigma]}$, until the first query answered "no". At this point, u' is extended to $u'\sigma$. If the length of $u'\sigma$ is now ℓ , then $u = u'\sigma$ is the desired counterexample; otherwise, we continue extending u'.

Note that if the input M is deterministic, then all of the restricted subset queries are made with deterministic finite acceptors. If L is recognized by a deterministic finite acceptor T_L , then the value of ℓ is bounded by $|M| \cdot |T_L|$, and the algorithm runs in time bounded by a polynomial in |M| and $|T_L|$.

We now turn to the ω -word case.

Theorem 6.3. There is an algorithm \mathbf{R}^{ω} with input M and restricted subset query access about L, (a language recognized by a DBW acceptor T_L) that correctly answers the unrestricted subset query with M about L. The algorithm $\mathbf{R}^{\omega}(M)$ runs in time bounded by a polynomial in |M| and $|T_L|$. If M is a DBW acceptor, then all the restricted subset queries will also be with DBW acceptors.

Algorithm 3 : $\mathbf{R}^{\omega}(M)$, implementing USQ(M)

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor **Ensure:** "yes" if $\llbracket M \rrbracket \subseteq L$, else "no" and (u, v) s.t. $u(v)^{\omega} \in (\llbracket M \rrbracket \setminus L)$ if RSQ(M) = "yes" then return "yes" else find $q \in F$ such that RSQ $(M_q) =$ "no" return "no" and Findctrex(M, q)

For the sake of generality, the proof considers subset queries with NBW acceptors. The procedure $\mathbf{R}^{\omega}(M)$ takes as input an NBW acceptor M, and has restricted subset query access (with NBW acceptors as inputs) to L; it is summarized in Algorithm 3. It first asks a restricted subset query with M about L, returning the answer "yes" if its query is answered "yes". Otherwise, for each $q \in F$, it constructs the acceptor $M_q = (Q, q_0, \delta, \{q\})$ with the single accepting state q and asks a restricted subset query with M_q about L, until the first query answered "no". There will be at least one such query answered "no" because any element of $(\llbracket M \rrbracket \setminus L)$ must visit at least one accepting state q of M infinitely many times, and will therefore be in $\llbracket M_q \rrbracket$. The procedure $\mathbf{R}^{\omega}(M)$ then calls the procedure Findctrex(M,q) to find a counterexample to return — i.e., a pair (u, v) such that $u(v)^{\omega} \in (\llbracket M_q \rrbracket \setminus L)$, and thus also $u(v)^{\omega} \in (\llbracket M \rrbracket \setminus L)$.

6.2. Producing a counterexample. The first challenge encountered in producing a counterexample, in comparison to the finite word case, is that one needs to work out both the period and the prefix of the counterexample to be found, and the two are correlated. Define $L_{q_0,q}$ to be the set of finite words that lead from the initial state q_0 to the state q in M, and define $L_{q,q}$ to be the set of nonempty finite words that lead from q back to q in M. Because the language $L_{q_0,q} \cdot (L_{q,q})^{\omega}$ is exactly the set of strings recognized by M_q , we know that $L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L \neq \emptyset$.

The procedure Findctrex(M, q) first finds a suitable period, corresponding to a bounded size of a prefix yet to be found, and then finds a prefix of that size in a similar manner to the finite word case. An example is shown in Appendix B.

Algorithm 4 : Findctrex(M,q)

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor; $q \in F$; $L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L \neq \emptyset$ **Ensure:** (u, v) such that $u(v)^{\omega} \in (L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L)$ let v = Findperiod(M, q)let u = Findprefix(M, q, v)**return** (u, v)

Since finding the period is more challenging than the prefix, we explain the procedure Findprefix(M, q, v) first. The procedure Findprefix(M, q, v), summarized in Algorithm 5, finds a prefix word u given a period word v which loops on state q and is guaranteed to be a period of a valid counterexample. It first finds a length k such that there exists $u \in L_{q_0,q}$ of length k such that $uv^{\omega} \notin L$. Then it finds such a word u symbol by symbol. Note that it uses length and prefix restricted versions of $L_{q_0,q}$.

Algorithm 5 : Findprefix(M, q, v)

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor; $q \in F$; $v \in L_{q,q}$; $L_{q_0,q} \cdot (v)^{\omega} \setminus L \neq \emptyset$ **Ensure:** $u \in L_{q_0,q}$ such that $u(v)^{\omega} \in (L_{q_0,q} \cdot (v)^{\omega} \setminus L)$ search for nonnegative integer k such that $\operatorname{RSQ}(L_{q_0,q}[k] \cdot (v)^{\omega}) =$ "no" let $u = \varepsilon$ **while** |u| < k **do** find $\sigma \in \Sigma$ such that $\operatorname{RSQ}(L_{q_0,q}[k, u\sigma] \cdot (v)^{\omega}) =$ "no" set $u = u \cdot \sigma$ **return** u

Finding the periodic part is much more challenging. Indeed, even if one knows that there is a period of the form $(a\Sigma^{\ell})^{\omega}$ for some ℓ then the size of the smallest period may be bigger than $\ell + 1$. For instance, if $L = \Sigma^{\omega} \setminus (abbaccadd)^{\omega}$ then there is a period of the form $(a\Sigma^2)^{\omega}$ but the shortest period of a counterexample is of size 9.

Procedure Findperiod(M, q), summarized in Algorithm 6, starts from the condition

$$L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L \neq \emptyset$$

and finds a sequence of words $v_1, v_2, \ldots \in L_{q,q}$ such that for each $n \ge 1$,

$$L_{q_0,q} \cdot (v_1 v_2 \cdots v_n \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

For a sufficiently long such sequence, there exists a subsequence $v = (v_i \cdots v_j)$ that is a suitable period word, as we prove in Section 7.2.

The procedure Nextword(M, q, y), summarized in Algorithm 7, is called with $y = v_1v_2\cdots v_n$ and finds a suitable next word v_{n+1} . After determining a length ℓ , it repeatedly

Algorithm 6 : Findperiod(M,q)

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor; $q \in F;$ $L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L \neq \emptyset$ **Ensure:** $v \in L_{q,q}$ such that $L_{q_0,q} \cdot (v)^{\omega} \setminus L \neq \emptyset$ let $y = \varepsilon$ for all integers $n = 1, 2, 3, \ldots$ do let $v_n = Nextword(M, q, y)$ set $y = y \cdot v_n$ for integers i, j with $1 \le i \le j \le n$ do for k = 0 to n|M| do if $\operatorname{RSQ}(L_{q_0,q}[k] \cdot (v_i \cdots v_j)^{\omega}) =$ "no" then return $v = v_i \cdots v_j$

calls the procedure Nextsymbol(M, q, y, ℓ, v') to determine the next symbol of a suitable word of length ℓ .

Algorithm 7 : Nextword(M, q, y)

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor; $q \in F;$ FIGURE 3. An illustration $y \in L_{q,q}$ or $y = \varepsilon$; of algorithm Nextword $L_{q_0,q} \cdot (y \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset$ **Ensure:** $v' \in L_{q,q}$ such that $L_{q_0,q} \cdot (y \cdot v' \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset$ $L_{q,q}$ search for integers $k, \ell \geq 0$ s.t. $\operatorname{RSQ}(L_{q_0,q}[k] \cdot (y \cdot L_{q,q}[\ell])^{\omega}) = \text{``no''}$ let $v' = \varepsilon$ while $|v'| < \ell$ do let $\sigma = Nextsymbol(M, q, y, \ell, v')$ $v_1 \in L$ set $v' = v' \cdot \sigma$ return v'

The procedure $Nextsymbol(M, q, y, \ell, v')$, summarized in Algorithm 8, is called to find a feasible next symbol with which to extend v' in the procedure Nextword.

Algorithm 8 : Nextsymbol(M, q, y, ℓ, v')

Require: RSQ access to L; $M = (Q, q_0, \delta, F)$, an NBW acceptor; $q \in F;$ $y \in L_{q,q};$ $v' \in \Sigma^*, |v'| < \ell;$ $L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell,v'] \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset$ **Ensure:** $\sigma \in \Sigma$ such that $L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell, v'\sigma] \cdot L_{q,q})^{\omega} \setminus L \neq$ Ø find integers $k \ge 0, m \ge 1$, and $\sigma \in \Sigma$ such that $\operatorname{RSQ}(L_{q_0,q}[k] \cdot (y \cdot L_{q,q}[\ell, v'\sigma] \cdot L_{q,q}[m])^{\omega}) = \text{``no''}$ return σ



 $v_3 \in L_{q,q}$

 $\in L_{q,q}$



7. Correctness

The main hurdle in proving the correctness of algorithm \mathbf{A}_{Trees} is to prove Theorem 6.3. The polynomial bound in the proof of Theorem 6.3 is obtained through a sequence of lemmas bounding the size of the acceptors used in \mathbf{A}_{Trees} subprocedures and the length restrictions and running time in calls to RSQ made by these procedures. Section 7.1 deals with bounding the acceptors, and Section 7.2 deals with the more challenging part, providing the length restrictions. Finally, Section 7.3 concludes with the theorem stating the correctness of algorithm \mathbf{A}_{Trees} .

7.1. Bounding the Acceptors. We turn to the representation (as NBW or DBW acceptors) of the languages used in restricted subset queries by $\mathbf{R}^{\omega}(M)$ and its subprocedures. We consider the size, out-degree, and time to construct the acceptors.

In $\mathbf{R}^{\omega}(M)$, there is a restricted subset query with M itself, and if that query is answered "no", a sequence of restricted subset queries with M_q for accepting states q until an answer of "no". Clearly, if M is an NBW acceptor, each M_q is an NBW acceptor of the same size and out-degree and is easily constructed from M, and similarly if M is a DBW acceptor.

The restricted subset queries made in *Findctrex* and its subprocedures are of the form $P \cdot (S)^{\omega}$, where P is a length and prefix restricted version of $L_{q_0,q}$ and S is a concatenation of (at most) a finite word and two length and prefix restricted versions of $L_{q,q}$. Therefore in what follows we consider the operations of concatenation and ω -repetition of regular languages of finite words.

These operations are particularly simple for DFW or NFW acceptors in *special form*, that is, containing at most one accepting state, which has no out-transitions defined. In general, any NFW acceptor can be converted to special form, possibly at the cost of increasing its out-degree. A regular language of finite words is recognized by a DFW acceptor in special form iff it is prefix-free.

However, if M is an NBW (resp., DBW) acceptor, then the finite word languages $L_{q_0,q}$ and $L_{q,q}$ are recognized by easily constructed NFW (resp., DFW) acceptors of size at most |M| and out-degree at most the out-degree of M. Lemma 6.2 shows that the length and prefix restricted versions of $L_{q_0,q}$ and $L_{q,q}$ are recognized by NFW (resp., DFW) acceptors in special form which may be constructed in time polynomial in |M|, ℓ , and |v| and have out-degree at most the out-degree of M.

Lemma 7.1. Suppose M_1 is an NFW acceptor in special form and M_2 is an NFW or NBW acceptor. Then an acceptor M for $[M_1] \cdot [M_2]$ can be constructed such that

(1)
$$|M| \leq |M_1| + |M_2|$$
,

- (2) the out-degree of M is at most the maximum of out-degrees of M_1 and M_2 ,
- (3) M can be constructed in polynomial time,
- (4) M is deterministic if M_1 and M_2 are deterministic,
- (5) M is an NFW in special form if M_2 is an NFW in special form.

Proof. Assume the states of M_1 and M_2 are disjoint. If M_1 has no accepting state then $[\![M_1]\!] = \emptyset$ and we take M to be a one-state acceptor of the same kind as M_2 that recognizes \emptyset . Otherwise, M_1 has one accepting state q_1 with no out transitions. If q_1 is also the initial state of M_1 , then $[\![M_1]\!] = \{\varepsilon\}$ and we take $M = M_2$.

Otherwise, M is constructed by taking the union of the two machines, removing the state q_1 and redirecting all the transitions to q_1 in M_1 to the initial state of M_2 . The initial

state of M is set to be the initial state of M_1 , and the accepting states of M are set to be the accepting states of M_2 .

Then M is an NFW acceptor if M_2 is an NFW acceptor, and an NBW acceptor if M_2 is an NBW acceptor. It is straightforward to verify the required properties of M.

Lemma 7.2. Suppose M_1 is an NFW acceptor in special form. Then an NBW acceptor M for $[\![M_1]\!]^{\omega}$ can be constructed such that

(1) $|M| \leq |M_1|$,

(2) the out-degree of M is at most the out-degree of M_1 ,

(3) M can be constructed in polynomial time,

(4) M is deterministic if M_1 is deterministic.

Proof. If M_1 has no accepting states then $\llbracket M_1 \rrbracket = \emptyset$. Otherwise, M_1 has one accepting state with no out transitions. If the accepting state of M_1 is also its initial state, then $\llbracket M_1 \rrbracket = \{\varepsilon\}$. In these two cases, $\llbracket M_1 \rrbracket^{\omega} = \emptyset$ and we take M to be an NBW acceptor with one state and no accepting states.

Otherwise, we construct M by removing from M_1 its unique accepting state q_1 and redirecting all the transitions into q_1 to the initial state of M_1 . The initial state of M_1 becomes the unique accepting state of M. It is straightforward to verify the required properties of M.

The above give us the following corollary for the procedure \mathbf{R}^{ω} .

Corollary 7.3. When the input to $\mathbf{R}^{\omega}(M)$ is an NBW (resp., DBW) acceptor M, each RSQ can be made with an NBW (resp., DBW) acceptor whose out-degree is at most the out-degree of M and can be constructed in time polynomial in |M| and parameters giving the length restrictions and the lengths of any words that appear.

7.2. Length restrictions and time bounds. We now turn to establish the correctness and running time of the subprocedures. The first two lemmas allow us to bound the parameters giving the length restrictions in inputs to RSQ.

Lemma 7.4. Let $S \subseteq L_{q,q}$ and suppose $L_{q_0,q} \cdot (S)^{\omega} \setminus L \neq \emptyset$. Then for some $k < |M| \cdot |T_L|$ we have $L_{q_0,q}[k] \cdot (S)^{\omega} \setminus L \neq \emptyset$.

Proof. Let $u = \sigma_1 \cdots \sigma_k$ be chosen to be a shortest word in $L_{q_0,q}$ such that $u \cdot (S)^{\omega} \setminus L \neq \emptyset$. Then for some s_1, s_2, \ldots from S, the ω -word

$$w = u \cdot s_1 \cdot s_2 \cdots$$

is in $(L_{q_0,q} \cdot (S)^{\omega} \setminus L)$.

There is an accepting run $r = r_0, r_1, \ldots$ of M on w. Let $t = t_0, t_1, \ldots$ be the unique run of the DBW acceptor T_L on w, which is rejecting. Consider the sequence of pairs (r_n, t_n) for $0 \le n \le |u|$. If $|u| \ge |M| \cdot |T_L|$, there will be a repeated pair, say $(r_i, t_i) = (r_j, t_j)$ for i < j. If we excise symbols i + 1 to j of u to get u' and the corresponding states from the runs rand t to get r' and t', we have

$$w' = u' \cdot s_1 \cdot s_2 \cdots$$

is accepted by M (witnessed by r') and rejected by T_L (witnessed by t'), so u' is a shorter word such that $u' \cdot (S)^{\omega} \setminus L \neq \emptyset$, a contradiction.

Lemma 7.5. Let $S \subseteq L_{q,q}$ and suppose $L_{q_0,q} \cdot (S \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset$. Then for some $k, \ell < \ell$ $|M| \cdot |T_L|$, we have that $L_{q_0,q}[k] \cdot (S \cdot L_{q,q}[\ell])^{\omega} \setminus L \neq \emptyset$.

Proof. Let $w \in (L_{q_0,q} \cdot (S \cdot L_{q,q})^{\omega} \setminus L)$. The unique run of the DBW acceptor T_L on w is rejecting, and does not visit an accepting state of T_L after some finite prefix. Because $S \subseteq L_{q,q}$, we may choose a sufficiently long prefix u of w such that $u \in L_{q_0,q}$ and when processing w, T_L never visits an accepting state after reading the prefix u.

Then w may be factored as

$$w = u(s_1 x_1)(s_2 x_2) \cdots,$$

where each $s_n \in S$ and each $x_n \in L_{q,q}$. There is an accepting run $r = r_0, r_1, \ldots$ of M on w, which we may assume visits the state q after u, and also after every s_n and every x_n .

Consider the states t_1, t_2, \ldots visited by T_L at the start of every group $(s_n x_n)$ when processing w. After at most $|T_L|$ groups, there must be a repeat, say $t_i = t_{i+p}$ for some p > 0. Let j = i + p - 1 and consider the ω -word

$$w' = u \cdot (s_1 x_1) \cdots (s_{i-1} x_{i-1}) \cdot ((s_i x_i) \cdots (s_j x_j))^{\omega}$$

There is an accepting run of M on w', and the unique run of T_L on w' is rejecting. Let

$$u' = u \cdot (s_1 x_1) \cdots (s_{i-1} x_{i-1})$$
 and $z = x_i \cdot (s_{i+1} x_{i+1}) \cdots (s_j x_j)$

Then $w' = u' \cdot (s_i z)^{\omega}$ and $u' \in L_{q_0,q}$ and $z \in L_{q,q}$. Consider an accepting run $r' = r'_0, r'_1, \ldots$ of M on w' that visits state q after processing u' and each occurrence of s_i and z. Consider the unique run $t = t'_0, t'_1, \ldots$ of T_L on w', which is rejecting. As in the proof of Lemma 7.4, if $|z| \ge |M| \cdot |T_L|$ then we may remove a segment of z that produces a cycle in the pairs (r'_n, t'_n) . Thus, for some $\ell < |M| \cdot |T_L|$, we have

$$L_{q_0,q} \cdot (S \cdot L_{q,q}[\ell])^{\omega} \setminus L \neq \emptyset.$$

Applying Lemma 7.4, there also exists $k < |M| \cdot |T_L|$ such that

$$L_{q_0,q}[k] \cdot (S \cdot L_{q,q}[\ell])^{\omega} \setminus L \neq \emptyset.$$

We now prove the correctness and polynomial running time of *Findprefix* and *Findperiod*, which establishes the correctness and polynomial running time of *Findctrex*.

Lemma 7.6. Assume $v \in L_{q,q}$ is such that

$$L_{q_0,q} \cdot (v)^{\omega} \setminus L \neq \emptyset.$$

Then in time polynomial in |M|, $|T_L|$ and |v|, the procedure Findprefix(M, q, v) returns a word $u \in L_{q_0,q}$ such that

$$u(v)^{\omega} \in (L_{q_0,q} \cdot (v)^{\omega} \setminus L)$$

Proof. The algorithm asks restricted subset queries about L for $\ell = 0, 1, 2, \ldots$ to find the least ℓ such that

$$L_{q_0,q}[\ell] \cdot (v)^{\omega} \setminus L \neq \emptyset$$

The value of ℓ is bounded by $|M| \cdot |T_L|$, by Lemma 7.4. It then searches symbol by symbol for a string u of length ℓ satisfying the required condition.

The procedure Findperiod depends on the procedures Nextword and Nextsymbol. The next lemma establishes the correctness and running time of the procedure Nextsymbol.

Lemma 7.7. Suppose ℓ is a positive integer, $y \in L_{q,q}$ or $y = \varepsilon$ and $v' \in \Sigma^*$ is such that $|v'| < \ell$ and we have

$$L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell, v'] \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

Then in time polynomial in |M|, $|T_L|$, |y| and ℓ , Nextsymbol (M, q, y, ℓ, v') finds a symbol $\sigma \in \Sigma$ such that

$$L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell, v'\sigma] \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

Proof. Consider an ω -word

$$w = u(yv'x_1y_1)(yv'x_2y_2)(yv'x_3y_3)\cdots,$$

in the language

$$L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell, v'] \cdot L_{q,q})^{\omega} \setminus L,$$

where $u \in L_{q_0,q}$, and for all $i, v'x_i \in L_{q,q}[\ell, v']$ and $y_i \in L_{q,q}$. Fix a particular accepting run of M_q on w that visits q after u and after every occurrence of $y, v'x_i$ and y_i in the factorization of w above.

Because in this run q is visited infinitely many times, we may assume that the prefix u is chosen so that w visits no accepting state of T_L after the prefix u has been processed. Now consider the sequence t_1, t_2, t_3, \ldots of states of T_L visited by w at the start of every group $(yv'x_iy_i)$. This sequence must repeat states of T_L , say $t_i = t_{i+p}$ for some p > 0. Let j = i + p - 1 and consider the word

$$w' = u(yv'x_1y_1)\cdots(yv'x_{i-1}y_{i-1})((yv'x_iy_i)\cdots(yv'x_jy_j))^{\omega}$$

Clearly, $w' \notin L$ because after the prefix u, w' visits only rejecting states of T_L .

Consider the cycle

$$((yv'x_iy_i)\cdots(yv'x_jy_j)).$$

If it is of length 1 (that is i = j), then we may duplicate the one group $(yv'x_iy_i)$ to make a cycle of length 2 without changing w'. Then we may factor the cycle as

$$((yv'x_iy_i)z)$$
 where $z = (yv'x_{i+1}y_{i+1})\cdots(yv'x_jy_j)$

and $z \in L_{q,q}$. Choosing σ to be the first symbol of x_i and x'_i to be the rest of x_i , we have

$$w' = u'(yv'\sigma x_i'z)^{\omega},$$

where $u' = u(yv'x_1y_1)\cdots(yv'x_{i-1}y_{i-1})$ and therefore

$$w' \in L_{q_0,q} \cdot (y \cdot L_{q,q}[\ell, v'\sigma] \cdot L_{q,q})^{\omega}$$

Thus we are guaranteed that some symbol σ with the required property exists. Lemma 7.5 (with $S = \{y\} \cdot L_{q,q}[\ell, v'\sigma]$) shows that there exist $k, m < |M| \cdot |T_L|$ such that

$$L_{q_0,q}[k] \cdot (y \cdot L_{q,q}[\ell, v'\sigma] \cdot L_{q,q}[m])^{\omega} \setminus L \neq \emptyset.$$

Thus, the search for k and m in the procedure Nextsymbol can enumerate such pairs (k, m) in increasing order of their maximum and try all $\sigma \in \Sigma$ for each pair until a suitable symbol σ is found to return. This process runs in time polynomial in |M|, $|T_L|$, |y| and ℓ .

Lemma 7.8. Suppose $y \in L_{q,q}$ or $y = \varepsilon$ is such that

$$L_{q_0,q} \cdot (y \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

Then in time bounded by a polynomial in |M|, $|T_L|$ and |y|, Nextword(M, q, y) returns a word $v' \in L_{q,q}$ of length bounded by $|M| \cdot |T_L|$ such that

$$L_{q_0,q} \cdot (yv' \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

Proof. By Lemma 7.5 (with $S = \{y\}$), the search for k and ℓ will succeed with both less than $|M| \cdot |T_L|$. Then ℓ calls to the procedure Nextsymbol will produce the required word v' of length ℓ .

The next lemma shows that Findperiod calls Nextword at most $|T_L|$ times.

Lemma 7.9. Suppose $v_1, v_2, \ldots, v_n \in L_{q,q}$ are such that

$$L_{q_0,q} \cdot (v_1 v_2 \cdots v_n \cdot L_{q,q})^{\omega} \setminus L \neq \emptyset.$$

Also suppose that the number of states of T_L is less than n. Then there exist integers i and j with $1 \le i \le j \le n$ such that

$$L_{q_0,q} \cdot (v_i v_{i+1} \cdots v_j)^{\omega} \setminus L \neq \emptyset$$

Proof. Consider an ω -word

$$w = u(v_1v_2\cdots v_n \cdot y_1)(v_1v_2\cdots v_n \cdot y_2)(v_1v_2\cdots v_n \cdot y_3)\cdots$$

in the language

$$L_{q_0,q} \cdot (v_1 v_2 \cdots v_n \cdot L_{q,q})^{\omega} \setminus L,$$

where $u \in L_{q_0,q}$ and each $y_i \in L_{q,q}$. Fix a particular accepting run of M on w in which state q is visited after each of the individual segments of w.

Considering the sequence of states of T_L that are visited in processing w, there must be some finite prefix after which only rejecting states of T_L are visited. Because the run of Mon w visits q infinitely often, we may assume that the prefix u of w extends past the last visit of T_L to an accepting state. Now consider the states t_1, t_2, \ldots, t_n visited by T_L at the start of each of the first occurrences of v_1, v_2, \ldots, v_n , respectively. Because n is greater than the number of states of T_L , some state of T_L must repeat in this sequence, say $t_i = t_{i+p}$ for some p > 0. Let j = i + p - 1 and consider the ω -word

$$w' = uv_1v_2\cdots v_{i-1}(v_iv_{i+1}\cdots v_j)^{\omega}.$$

Then $w' \in L_{q_0,q} \cdot (v_i v_{i+1} \cdots v_j)^{\omega}$ because $u' = u v_1 v_2 \cdots v_{i-1}$ is in $L_{q_0,q}$. However, because only rejecting states of T_L are visited in the repeating portion of the word, $w' \notin L$.

The final lemma, presented below, establishes the correctness and polynomial running time of the procedure *Findperiod*.

Lemma 7.10. Suppose $L_{q_0,q} \cdot (L_{q,q})^{\omega} \setminus L \neq \emptyset$. Then, in polynomial time in |M| and $|T_L|$, the procedure Findperiod(M,q) with restricted query access to L returns a period word v satisfying the condition

$$L_{q_0,q} \cdot (v)^{\omega} \setminus L \neq \emptyset.$$

Proof. The preconditions of Findperiod are satisfied, and it calls Nextword(M, q, y) repeatedly, with $y = \varepsilon$, then $y = v_1$, then $y = v_1v_2$, and so on, where v_{n+1} is the value returned by the call with $y = v_1v_2 \cdots v_n$. Each of these calls satisfies the preconditions of Nextword, so after at most $|T_L|$ such calls, Findperiod returns a correct period word v, by Lemma 7.9. \Box

These lemmas can be used in combination to prove Theorem 6.3, giving a polynomial time reduction of unrestricted subset queries to restricted subset queries for NBW acceptors (resp., DBW acceptors.)

7.3. Correctness of $\mathbf{A}_{\text{Trees}}$. The lemmas established in the previous subsection also show the correctness and running time of Findctrex(M',q) when called by $\mathbf{A}_{\text{Trees}}$, provided that each RSQ about L is correctly answered and q satisfies the precondition of Findctrex.

To complete the consideration of representation issues, we must prove that \mathbf{A}_{Trees} can successfully simulate *Findctrex* as stated in Lemma 7.11.

Lemma 7.11. When $\mathbf{A}_{\text{Trees}}$ simulates Findctrex(M', q) in response to a negative counterexample t, every RSQ can be simulated with a MQ about $\text{Trees}_d(L)$.

Proof. In the learning algorithm \mathbf{A}_{Trees} , when a negative counterexample t represented by A_t is received, the algorithm simulates the procedure Findctrex(M', q) where $M' = acceptor(A_t)$ is a NBW acceptor recognizing paths(t) and q is an accepting state of M'. Note that by Lemma 5.1, $|M'| \leq |A_t|$ and the out-degree of M' is at most d, the arity of t.

Then Corollary 7.3 shows that each RSQ is with a NBW acceptor that has out-degree at most the out-degree of M', which is at most d. Also, each such NBW acceptor can be constructed in time polynomial in |M'| and parameters giving the length restrictions and the lengths of any words that appear.

The final observation is that each such RSQ is made with an NBW acceptor that recognizes a safety language of the form $P \cdot (S)^{\omega}$, where P and S are each languages of fixed-length finite words. Then, by Lemma 5.2 each such RSQ(N) can be simulated by \mathbf{A}_{Trees} using $MQ(tree_d(N))$ about $Trees_d(L)$.

If q does not satisfy the precondition of *Findctrex*, then the procedure may run forever. However, at least one accepting state q satisfies the precondition, so at least one simulation will halt and return (u, v), at which point \mathbf{A}_{Trees} terminates all the other simulations. This concludes the proof of the reduction given by \mathbf{A}_{Trees} , whose general statement is given in Theorem 7.12 below.

Theorem 7.12. Suppose $\mathbb{C} \subseteq \mathbb{DBW}$ and \mathbf{A} is a polynomial time algorithm that learns class \mathbb{C} using membership and equivalence queries. Then for every positive integer d there is a polynomial time algorithm $\mathbf{A}_{\text{Trees}}$ that learns the class $\text{Trees}_d(\mathbb{C})$ using membership and equivalence queries.

This theorem, together with Maler and Pnueli's [MP95] polynomial time algorithm to learn the class of weak regular ω -word languages using membership and equivalence queries proves our main result — Theorem 4.1.

8. DISCUSSION

We have shown that if $\mathbb{C} \subseteq \mathbb{DBW}$ can be learned in polynomial time with membership and equivalence queries, then $Trees_d(\mathbb{C})$ can be learned in polynomial time with membership and equivalence queries for all $d \geq 1$. Consequently, there is a polynomial time algorithm to learn $Trees_d(\mathbb{DwPW})$ with membership and equivalence queries. We have also shown that there are polynomial time algorithms that implement unrestricted subset queries using restricted subset queries for \mathbb{DFW} , \mathbb{NFW} , \mathbb{DBW} and \mathbb{NBW} .

One open question is whether there is an interesting subclass of \mathbb{DBW} that is larger than \mathbb{DwPW} but still learnable in polynomial time using membership and equivalence queries, to which Theorem 7.12 would also apply.

Acknowledgment

The authors would like to thank the anonymous reviewers for their valuable feedback and helpful suggestions. This research was supported by the United States - Israel Binational Science Foundation, Jerusalem, Israel (BSF) under grant number #8758451 and by the Office of Naval Research (ONR) under grant number #N00014-17-1-2787.

References

- [AdRP13] F. Aarts, J. de Ruiter, and E. Poll. Formal models of bank cards for free. In Sixth IEEE International Conference on Software Testing, Verification and Validation, ICST 2013 Workshops Proceedings, Luxembourg, Luxembourg, March 18-22, 2013, pages 461–468, 2013.
- [AF16] D. Angluin and D. Fisman. Learning regular omega languages. Theor. Comput. Sci., 650:57–72, 2016.
- [Ang87] D. Angluin. Learning regular sets from queries and counterexamples. Inf. Comput., 75(2):87–106, 1987.
- [Ang88] D. Angluin. Queries and concept learning. Mach. Learn., 2(4):319–342, 1988.
- [Ang90] D. Angluin. Negative results for equivalence queries. Mach. Learn., 5(2):121–150, 1990.
- [Ang16] D. Angluin. The class DBW \cap DCW of ω -languages is identifiable in the limit from positive data and membership queries with polynomial time and data. Technical Report YALEU/DCS/TR-1528, Dept. of Computer Science, Yale University, August 2016.
- [BKKS13] U. Boker, D. Kuperberg, O. Kupferman, and M. Skrzypczak. Nondeterminism in the presence of a diverse or unknown future. In Automata, Languages, and Programming, 40th Inter. Colloq., ICALP, pages 89–100, 2013.
- [CBSS10] C. Y. Cho, D. Babic, E. C. R. Shin, and D. Song. Inference and analysis of formal models of botnet command and control protocols. In *Proceedings of the 17th ACM Conference on Computer* and Communications Security, CCS 2010, Chicago, Illinois, USA, October 4-8, 2010, pages 426–439, 2010.
- [CGP03] J. M. Cobleigh, D. Giannakopoulou, and C. S. Păsăreanu. Learning assumptions for compositional verification. In Proc. of the 9th Inter. Conf. on Tools and Algorithms for the Construction and Analysis of Systems, TACAS, pages 331–346, 2003.
- [CNP94] H. Calbrix, M. Nivat, and A. Podelski. Ultimately periodic words of rational ω-languages. In Proc. of the 9th Inter. Conf. on Mathematical Foundations of Programming Semantics, pages 554–566, 1994.
- [CPPdR14] G. Chalupar, S. Peherstorfer, E. Poll, and J. de Ruiter. Automated reverse engineering using Lego®. In 8th USENIX Workshop on Offensive Technologies, WOOT '14, San Diego, CA, USA, August 19, 2014., 2014.
- [dRP15] J. de Ruiter and E. Poll. Protocol state fuzzing of TLS implementations. In 24th USENIX Security Symposium, USENIX Security 15, Washington, D.C., USA, August 12-14, 2015., pages 193–206, 2015.
- [ES88] E. A. Emerson and J. Srinivasan. Branching time temporal logic. In Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, School/Workshop, Noordwijkerhout, The Netherlands, May 30 - June 3, 1988, Proceedings, pages 123–172, 1988.
- [FCC⁺08] A. Farzan, Y. Chen, E.M. Clarke, Y. Tsay, and B. Wang. Extending automated compositional verification to the full class of omega-regular languages. In *Tools and Algorithms for the Construction and Analysis of Systems, TACAS*, volume 4963 of *LNCS*, pages 2–17, 2008.
- [FLM⁺13] L. Feng, S. Lundmark, K. Meinke, F. Niu, M. A. Sindhu, and P. Y. H. Wong. Case studies in learning-based testing. In *ICTSS*, volume 8254 of *Lecture Notes in Computer Science*, pages 164–179. Springer, 2013.
- [GTW02] E. Grädel, W. Thomas, and T. Wilke, editors. Automata, Logics, and Infinite Games: A Guide to Current Research. Springer, 2002.
- [HHNS02] A. Hagerer, H. Hungar, O. Niese, and B. Steffen. Model generation by moderated regular extrapolation. In FASE, volume 2306 of Lecture Notes in Computer Science, pages 80–95. Springer, 2002.

M. Jayasrirani, M. Humrosia Begam, and D. Gnanaraj Thomas. Learning of regular omega-tree

	languages. In Grammatical Inference: Algorithms and Applications, 9th Inter. Colloq., ICGI, pages 295–297, 2008.
[KS15]	D. Kuperberg and M. Skrzypczak. On determinisation of good-for-games automata. In Automata, Languages and Programming 12nd Inter Collog ICALP pages 299-310 2015
[KSV06]	 D. Kupferman, S. Safra, and M. Y. Vardi. Relating word and tree automata. Annals of Pure and Annlied Logic, 138:126–146, 2006.
[KVW00]	O. Kupferman, M. Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. J. ACM, 47(2):312–360, 2000.
[Löd11]	C. Löding. Automata on infinite trees. Available online from author, 2011.
[MNRS04]	T. Margaria, O. Niese, H. Raffelt, and B. Steffen. Efficient test-based model generation for legacy reactive systems. In <i>HLDVT</i> , pages 95–100. IEEE Computer Society, 2004.
[MP90]	Z. Manna and A. Pnueli. A hierarchy of temporal properties. In <i>Proc. of the Ninth Annual ACM Symposium on Principles of Distributed Computing</i> , <i>PODC</i> , pages 377–410, 1990.
[MP95]	O. Maler and A. Pnueli. On the learnability of infinitary regular sets. <i>Inform. Comput.</i> , 118(2):316–326, 1995.
[NA06]	W. Nam and R. Alur. Learning-based symbolic assume-guarantee reasoning with automatic decomposition. In <i>ATVA</i> , volume 4218 of <i>Lecture Notes in Computer Science</i> , pages 170–185. Springer 2006
[NW98]	 D. Niwiński and I. Walukiewicz. Relating hierarchies of word and tree automata. In In Symposium on Theoretical Aspects in Computer Science, STACS, LNCS 1373, pages 320–331. Springer-Verlag, 1008
[PP04]	D. Perrin and JÉ. Pin. Infinite Words: Automata, Semigroups, Logic and Games. Elsevier, 2004.
[PVY02]	D. A. Peled, M. Y. Vardi, and M. Yannakakis. Black box checking. <i>Journal of Automata</i> , Languages and Combinatorics, 7(2):225–246, 2002.
[Rab72]	M. Rabin. Automata on infinite objects and Church's problem. Regional Conf. series in mathe- matics. Conf. Board of the Math. Sci. by the American Mathematical Society, 1972.
[SG14]	M. Shahbaz and R. Groz. Analysis and testing of black-box component-based systems by inferring partial models. <i>Softw. Test.</i> , <i>Verif. Reliab.</i> , 24(4):253–288, 2014.
[SHV16]	M. Schuts, J. Hooman, and F. W. Vaandrager. Refactoring of legacy software using model learning and equivalence checking: An industrial experience report. In <i>IFM</i> , volume 9681 of <i>Lecture Notes in Computer Science</i> , pages 311–325. Springer, 2016.
[Vaa17]	F. Vaandrager. Model learning. Commun. ACM, 60(2):86–95, 2017.
[Wag75]	K. W. Wagner. A hierarchy of regular sequence sets. In 4th Symposium on Mathematical Foundations of Computer (MFCS), pages 445–449, 1975.
[WNS ⁺ 13]	S. Windmüller, J. Neubauer, B. Steffen, F. Howar, and O. Bauer. Active continuous quality control. In <i>CBSE</i> , pages 111–120. ACM, 2013.

Appendix A. Example of \mathbf{A}_{Trees}

We illustrate the algorithm \mathbf{A}_{Trees} learning the language $Trees_2(L)$ where L is the language recognized by the DBW pictured in Fig. 5. Note that L has the rejecting SCC {2}, which is a subset of the accepting SCC {1, 2}, so L is not accepted by any DCW, and is therefore not in $\mathbb{DBW} \cap \mathbb{DCW}$. We assume that the learning algorithm \mathbf{A}_{Trees} has MQ and EQ access to the language $Trees_2(L)$. \mathbf{A}_{Trees} also has access to an oracle \mathbf{A} that makes MQs and EQs about L and ultimately outputs a DBW recognizing L. The treatment of MQs is straightforward, so we focus on Eqs. To help illustrate the behavior of \mathbf{A}_{Trees} , we choose two hypothetical EQs that \mathbf{A} could make to L, as well as the possible counterexample trees to the resulting EQs made by \mathbf{A}_{Trees} .

Suppose the first EQ that **A** makes to *L* is with the DBW H_1 pictured in Fig. 5. The language recognized by H_1 is $L_1 = (a+b)(a+b+c)^{\omega}$, which is incomparable with *L*. The

[JBT08]



FIGURE 5. Left, the DBW M recognizing L. Right, H_1 , an EQ made by A to L.



FIGURE 6. Left, the regular ω -tree counterexample T_1 to $H_1^{T,2}$. Right, the top three levels of the extensive form of T_1 .

algorithm \mathbf{A}_{Trees} constructs the deterministic tree acceptor $H_1^{T,2}$, which recognizes all binary trees all of whose infinite paths are in L_1 , and makes an EQ to $Trees_2(L)$ with $H_1^{T,2}$. Suppose that the counterexample returned is the regular ω -tree T_1 pictured in Fig. 6, with the top three levels of the extensive form of T_1 also shown.

At this point, the $\mathbf{A}_{\text{Trees}}$ algorithm must call on the Accepted? procedure with inputs T_1 and H_1 to decide whether the counterexample T_1 is accepted or rejected by the hypothesis $H_1^{T,2}$. This procedure constructs the product automaton Π_1 shown in Fig. 7. Because Π_1 accepts $\{1,2\}^{\omega}$, the Accepted? procedure reports that $H_1^{T,2}$ accepts the ω -tree T_1 . Because T_1 is incorrectly accepted by $H_1^{T,2}$, the learning algorithm $\mathbf{A}_{\text{Trees}}$ constructs a DBW M_1^{ω} accepted $(T_1)^{\omega}$.

Because T_1 is incorrectly accepted by $H_1^{1,2}$, the learning algorithm \mathbf{A}_{Trees} constructs a DBW $M' = acceptor(T_1)$ accepting precisely all the paths of T_1 . The DBW M' is shown in Fig. 7. Because at least one ω -word accepted by M' must not be in the target language L, the procedure *Findctrex* is called with the DBW M' and restricted subset query access to the target language L. The restricted subset queries are simulated using MQs to $Trees_2(L)$ and the representation of a safety language as a regular ω -tree (Lemma 5.2).

Assume that Findctrex returns the pair (a, b), representing the ω -word ab^{ω} , which is accepted by H_1 and is not in L. At this point, the EQ made by algorithm **A** with the DBW H_1 can be answered with the pair (a, b).

Assume that at some later point, **A** makes an EQ to L with the DBW H_2 shown in Fig. 8. (Note that the language recognized by H_2 is $(a + ba)^{\omega}$, which is a proper subset of L.) **A**_{Trees} then makes an EQ to $Trees_2(L)$ with the ω -tree automaton $H_2^{T,2}$. Assume that the counterexample returned is the regular ω -tree T_2 , shown in Fig. 8. (It can be verified that the tree T_2 is in $Trees_2(L)$.)



FIGURE 7. Left, the product automaton Π_1 defined using T_1 and H_1 . Right, the DBW $M' = acceptor(T_1)$ recognizing all the paths in T_1 .



FIGURE 8. Left, the DBW H_2 , an EQ made by **A** to *L*. Right, the counterexample regular ω -tree T_2 .



FIGURE 9. The product DBW Π_2 defined using T_2 and H_2 .

Then \mathbf{A}_{Trees} calls Accepted? with the tree T_2 and the DBW H_2 . The Accepted? procedure constructs the product DBW Π_2 shown in Fig. 9. The DBW Π_2 does not accept all ω -words in $\{1,2\}^{\omega}$, for example, the ω -word $22(21)^{\omega}$ is not accepted. The corresponding input ω -word is $bc(ba)^{\omega}$, which is in L but is not accepted by H_2 . Thus, the procedure Accepted? could return the pair (bc, ba), which would then be supplied to \mathbf{A} as a counterexample to the EQ with H_2 .

APPENDIX B. EXAMPLE OF Findctrex

Consider the two ω -languages

$$[M_1] = a^* b((a+c)(b+c))^{\omega},$$

recognized by the DBW M_1 pictured in Fig. 10, and

 $[\![M_2]\!] = (a+c)^* b((a+c)a^*b)^{\omega},$

21:28



FIGURE 10. Left, the DBW M_1 . Right, the DBW M_2 .

recognized by the DBW M_2 pictured in Fig. 10. Note that $\llbracket M_2 \rrbracket$ is not a subset of $\llbracket M_1 \rrbracket$. For example, the ω -words $cb(ab)^{\omega}$ and $b(aab)^{\omega}$ are both in $(\llbracket M_2 \rrbracket \setminus \llbracket M_1 \rrbracket)$.

Assume that the procedure Findctrex is called with a restricted subset query oracle for $\llbracket M_1 \rrbracket$ and inputs consisting of the DBW M_2 and final state 2. It calls the procedure Findperiod with inputs M_2 and state 2 to get a period v, for example v = aab, with the property that some prefix followed by $(aab)^{\omega}$ is accepted by M_2 and is not in $\llbracket M_1 \rrbracket$. It then calls the procedure Findprefix with inputs M_2 , state 2 and the period aab to find a prefix ufor which $u(aab)^{\omega}$ is accepted by M_2 and is not in $\llbracket M_1 \rrbracket$, for example u = b, and returns the pair (b, aab) representing the ω -word $b(aab)^{\omega}$ accepted by M_2 but not in $\llbracket M_1 \rrbracket$.

In the Findperiod computation on inputs M_2 and state 2, the Nextword procedure is repeatedly called with inputs M_2 , state 2 and the word $v_1v_2\cdots v_{n-1}$ to find the next word v_n , until a restricted subset query yields a period $v = v_i \cdots v_j$ with the property that some prefix followed by $(v)^{\omega}$ is accepted by M_2 and is not in the language $[M_1]$. Findperiod returns the period v.

For M_2 we have the following.

$$\llbracket M_2 \rrbracket_{1,2} = (a+c)^* b((a+c)a^*b)^*$$
 and $\llbracket M_2 \rrbracket_{2,2} = ((a+c)a^*b)^+.$

We consider also the following length-restricted versions of these languages.

$$\begin{split} & [\![M_2]\!]_{1,2}[1] = b \\ & [\![M_2]\!]_{1,2}[2] = (a+c)b \\ & [\![M_2]\!]_{1,2}[3] = (a+c)(a+c)b + b(a+c)b \\ & [\![M_2]\!]_{2,2}[1] = \emptyset \\ & [\![M_2]\!]_{2,2}[2] = (a+c)b \\ & [\![M_2]\!]_{2,2}[3] = (a+c)ab. \end{split}$$

The computation of Nextword on inputs M_2 , state 2 and the initial value $y = \varepsilon$ first searches using restricted subset queries to find nonnegative integers k and ℓ such that there exist a prefix of length k in $[M_2]_{1,2}$ and a period of length ℓ in $[M_2]_{2,2}$ that yield an ω -word accepted by M_2 that is not in $[M_1]$. The value of ℓ is fixed, and a word v' of length ℓ is built up symbol by symbol calling Nextsymbol to yield the result of Nextword. In our example, k = 1 and $\ell = 2$ do not suffice (because $b((a + c)b)^{\omega}$ is a subset of $[[M_1]]$) but k = 2 and $\ell = 2$ do (because $(a + c)b((a + c)b)^{\omega}$ is not a subset of $[[M_1]]$) and k = 1 and $\ell = 3$ also do (because $b((a + c)ab)^{\omega}$ is not a subset of $[[M_1]]$). In this example, when ab, cb, aab or cab is returned by Nextword to Findperiod, the value v_1 suffices as the value of v to be returned by Findperiod. In more complex cases, repeated calls to Nextword may be necessary.