

COMPLEXITY OF MODEL CHECKING RECURSION SCHEMES FOR FRAGMENTS OF THE MODAL MU-CALCULUS

NAOKI KOBAYASHI AND C.-H. LUKE ONG

Graduate School of Information Sciences, Tohoku University, 6-3-09 Aoba, Aramaki, Aoba-ku
Sendai, 980-8579 Japan

URL: `koba@ecei.tohoku.ac.jp`

Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford
OX1 3QD, UK

URL: `Luke.Ong@cs.ox.ac.uk`

ABSTRACT. Ong has shown that the modal mu-calculus model checking problem (equivalently, the alternating parity tree automaton (APT) acceptance problem) of possibly-infinite ranked trees generated by order- n recursion schemes is n -EXPTIME complete. We consider two subclasses of APT and investigate the complexity of the respective acceptance problems. The main results are that, for APT with a single priority, the problem is still n -EXPTIME complete; whereas, for APT with a disjunctive transition function, the problem is $(n - 1)$ -EXPTIME complete. This study was motivated by Kobayashi's recent work showing that the resource usage verification of functional programs can be reduced to the model checking of recursion schemes. As an application, we show that the resource usage verification problem is $(n - 1)$ -EXPTIME complete.

1. INTRODUCTION

The model checking problem for higher-order recursion schemes has been a topic of active research in recent years (for motivation as to why the problem is interesting, see e.g. the introduction of Ong's paper [15]). This paper studies the complexity of the problem with respect to certain fragments of the modal μ -calculus. A higher-order recursion scheme (recursion scheme, for short) is a kind of (deterministic) grammar for generating a possibly-infinite ranked tree. The model checking problem for recursion schemes is to decide, given an order- n recursion scheme \mathcal{G} and a specification ψ for infinite trees, whether the tree generated by \mathcal{G} satisfies ψ . Ong [15] has shown that if ψ is a modal μ -calculus formula (or equivalently, an alternating parity tree automaton), then the model checking problem is n -EXPTIME complete.

Following Ong's work, Kobayashi [12] has recently applied the decidability result to the model checking of higher-order functional programs (precisely, programs of the simply-typed λ -calculus with recursion and resource creation/access primitives). He considered the

1998 ACM Subject Classification: F.3.1, D.2.4.

Key words and phrases: model checking, higher-order recursion schemes, modal mu-calculus, complexity.

resource usage verification problem [7]—the problem of whether programs access dynamically created resources in a valid manner (e.g. whether every opened file will eventually be closed, and thereafter never read from or written to before it is reopened). He showed that the resource usage verification problem reduces to a model checking problem for recursion schemes by giving a transformation that, given a functional program, constructs a recursion scheme that generates all possible resource access sequences of the program. From Ong’s result, it follows that the resource usage verification problem is in n -EXPTIME (where, roughly, n is the highest order of types in the program). This result also implies that various other verification problems, including (the precise verification of) reachability (“Does a closed program reach the fail command?”) and flow analysis (“Does a sub-term e evaluate to a value generated at program point l ?”), are also in n -EXPTIME, as they can be easily recast as resource usage verification problems.

It was however unknown whether n -EXPTIME is the tightest upper-bound of the resource usage verification problem. Although the model checking of recursion schemes is n -EXPTIME-hard for the full modal μ -calculus, only a certain fragment of the modal μ -calculus is used in Kobayashi’s approach to the resource usage verification problem. First, specifications are restricted to safety properties, which can be described by Büchi tree automata with a trivial acceptance condition (the class called “trivial automata” by Aehlig [1]). Secondly, specifications are also restricted to linear-time properties—the branching structure of trees is ignored, and only the path languages of trees are of interest. Thus, one may reasonably hope that there is a more tractable model checking algorithm than the n -EXPTIME algorithm.

The goal of this paper is, therefore, to study the complexity of the model checking of recursion schemes for various fragments of the modal μ -calculus (or, alternating parity tree automata) and to apply the result to obtain tighter bounds of the complexity of the resource usage verification problem.

The main results of this paper are as follows:

(i) The problem of whether a given Büchi tree automaton with a trivial acceptance condition (or, equivalently, alternating parity tree automaton with a single priority 0) accepts the tree generated by an order- n recursion scheme is still n -EXPTIME-hard, both in the size of the recursion scheme and that of the automaton. This follows from the n -EXPTIME-completeness of the word acceptance problem of higher-order alternating push-down automata¹ [6].

(ii) We introduce a new subclass of alternating parity tree automata (APT) called *disjunctive APT*, and show that its acceptance problem for trees generated by order- n recursion schemes is $(n - 1)$ -EXPTIME complete. From this general result, it follows that both the linear-time properties (including reachability, which is actually $(n - 1)$ -EXPTIME-complete) and finiteness of the tree generated by a recursion scheme are $(n - 1)$ -EXPTIME.

(iii) As an application, we show that the resource usage verification problem [12] is also $(n - 1)$ -EXPTIME-complete, where n is the highest order of types used in the source program (written in an appropriate language [12]).

The rest of this section is organized as follows. Section 2 reviews definitions of recursion schemes and alternating parity tree automata (APT). Section 3 introduces the class of trivial APT and studies the complexity of model checking recursion schemes. Section 4 introduces the class of disjunctive APT and studies the complexity of model checking recursion schemes.

¹Engelfriet’s proof [6] is for a somewhat different (but equivalent) machine which is called *iterated push-down automaton*.

Section 5 applies the result to analyze the complexity of the resource usage verification. Section 6 discusses related work and concludes the paper.

2. PRELIMINARIES

Let Σ be a ranked alphabet, i.e. a function that maps a terminal symbol to its arity, which is a non-negative integer. Let $\mathbb{N} = \{1, 2, \dots\}$. A Σ -labeled (unranked) tree T is a partial map from \mathbb{N}^* to $\text{dom}(\Sigma)$, such that $sk \in \text{dom}(T)$ (where $s \in \mathbb{N}^*$, $k \in \mathbb{N}$) implies $\{s\} \cup \{sj \mid 1 \leq j < k\} \subseteq \text{dom}(T)$. A (possibly infinite) sequence π over \mathbb{N} is a *path* of T just if every finite prefix of π is in $\text{dom}(T)$. A tree is *ranked* just if $\max\{j \mid sj \in \text{dom}(T)\}$ is equal to the arity of $T(s)$ for each $s \in \text{dom}(T)$.

Higher-Order Recursion Schemes. The set of *types* is defined by:

$$\kappa ::= \circ \mid \kappa_1 \rightarrow \kappa_2$$

where \circ is the type of trees. By convention, \rightarrow associates to the right; thus, for example, $o \rightarrow o \rightarrow o$ means $o \rightarrow (o \rightarrow o)$. The *order* of κ , written $\text{order}(\kappa)$, is defined by:

$$\begin{aligned} \text{order}(\circ) &:= 0 \\ \text{order}(\kappa_1 \rightarrow \kappa_2) &:= \max(\text{order}(\kappa_1) + 1, \text{order}(\kappa_2)). \end{aligned}$$

A (deterministic) *higher-order recursion scheme* (recursion scheme, for short) is a quadruple $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$, where

- (i) Σ is a ranked alphabet of *terminal symbols*.
- (ii) \mathcal{N} is a map from a finite set of symbols called *non-terminals* to types.
- (iii) \mathcal{R} is a set of rewrite rules $F \tilde{x} \rightarrow t$. Here $\tilde{x} = x_1, \dots, x_n$ abbreviates a sequence of variables, and t is an applicative term constructed from non-terminals, terminals, and variables x_1, \dots, x_n .
- (iv) S is a *start symbol*.

We require that $\mathcal{N}(S) = \circ$. The set of (typed) terms is defined in the standard manner: A non-terminal or variable of type κ is a term of type κ . A terminal of arity k is a term of type $\underbrace{\circ \rightarrow \dots \rightarrow \circ}_k \rightarrow \circ$. If terms t_1 and t_2 have types $\kappa_1 \rightarrow \kappa_2$ and κ_1 respectively, then $t_1 t_2$

is a term of type κ_2 . By convention, application associates to the left; thus, for example, $st u$ means $(st)u$. For each rule $F \tilde{x} \rightarrow t$, $F \tilde{x}$ and t must be terms of type \circ . There must be exactly one rewrite rule for each non-terminal. The *order* of a recursion scheme is the highest order of (the types of) its non-terminals.

A rewrite relation on terms is defined inductively by:

- (i) If $F \tilde{x} \rightarrow t \in \mathcal{R}$, then $F \tilde{s} \rightarrow_{\mathcal{G}} [\tilde{s}/\tilde{x}]t$.
- (ii) If $t \rightarrow_{\mathcal{G}} t'$, then $ts \rightarrow_{\mathcal{G}} t's$ and $st \rightarrow_{\mathcal{G}} st'$.

The *value tree* of a recursion scheme \mathcal{G} , written $\llbracket \mathcal{G} \rrbracket$, is the (possibly infinite) tree obtained by infinite rewriting of the start symbol S . More precisely, let us define t^\perp by:

$$a^\perp := a \quad F^\perp := \perp \quad (t_1 t_2)^\perp := \begin{cases} \perp & \text{if } t_1^\perp = \perp \\ t_1^\perp t_2^\perp & \text{otherwise} \end{cases}$$

The value tree $\llbracket \mathcal{G} \rrbracket$ is the $\Sigma \cup \{\perp \mapsto 0\}$ -ranked tree defined by:

$$\llbracket \mathcal{G} \rrbracket := \bigsqcup \{t^\perp \mid S \rightarrow_{\mathcal{G}}^* t\}.$$




Figure 1: The tree generated by the recursion scheme of Example 1

Here, $\sqcup S$ denotes the least upper bound with respect to the tree order \sqsubseteq defined by

$$T_1 \sqsubseteq T_2 \iff \forall s \in \text{dom}(T_1). (T_1(s) = T_2(s) \vee T_1(s) = \perp)$$

Note that $\llbracket \mathcal{G} \rrbracket$ is always well-defined, as the rewrite relation $\rightarrow_{\mathcal{G}}$ is confluent.

Example 1. Consider the recursion scheme $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$ where

$$\begin{aligned} \Sigma &= \{a \mapsto 2, b \mapsto 1, c \mapsto 1, e \mapsto 0\} \\ \mathcal{N} &= \{S \mapsto o, F \mapsto (o \rightarrow o) \rightarrow o \rightarrow o, I \mapsto o \rightarrow o, C \mapsto (o \rightarrow o) \rightarrow (o \rightarrow o) \rightarrow (o \rightarrow o)\} \\ \mathcal{R} &= \{ \\ &\quad S \rightarrow F I e, \\ &\quad F f x \rightarrow a(f x) (F (C b f) (c x)), \\ &\quad I x \rightarrow x, \\ &\quad C f g x \rightarrow f(g x) \\ &\quad \} \end{aligned}$$

S is reduced as follows.

$$\begin{aligned} S &\longrightarrow F I e \\ &\longrightarrow a(I e) (F (C b I) (c e)) \\ &\longrightarrow a e (a (C b I (c e)) (F (C b (C b I)))) (c (c e)) \\ &\longrightarrow^* a e (a (b (c e)) (F (C b (C b I)))) (c (c e)) \\ &\longrightarrow^* a e (a (b (c e)) (a (b^2 (c^2 e)) (a (b^3 (c^3 e)) \dots))) \end{aligned}$$

The value tree is shown in Figure 2. Each path of the tree is labelled by $a^{m+1}b^m c^m e$. \square

Alternating parity tree automata. Given a finite set X , the set $B^+(X)$ of *positive Boolean formulas* over X is defined as follows. We let θ range over $B^+(X)$.

$$\theta ::= t \mid f \mid x \mid \theta \wedge \theta \mid \theta \vee \theta$$

where x ranges over X . We say that a subset Y of X *satisfies* θ just if assigning true to elements in Y and false to elements in $X \setminus Y$ makes θ true.

An *alternating parity tree automaton* (or APT for short) over Σ -labelled trees is a tuple $\mathcal{A} = (\Sigma, Q, \delta, q_I, \Omega)$ where

- (i) Σ is a ranked alphabet; let m be the largest arity of the terminal symbols;
- (ii) Q is a finite set of states, and $q_I \in Q$ is the initial state;
- (iii) $\delta : Q \times \Sigma \longrightarrow \mathbf{B}^+(\{1, \dots, m\} \times Q)$ is the transition function where, for each $f \in \Sigma$ and $q \in Q$, we have $\delta(q, f) \in \mathbf{B}^+(\{1, \dots, \text{arity}(f)\} \times Q)$; and
- (iv) $\Omega : Q \longrightarrow \{0, \dots, M-1\}$ is the priority function.

A *run-tree* of an APT \mathcal{A} over a Σ -labelled ranked tree T is a $(\text{dom}(T) \times Q)$ -labelled unranked tree r satisfying:

- (i) $\epsilon \in \text{dom}(r)$ and $r(\epsilon) = (\epsilon, q_I)$; and
- (ii) for every $\beta \in \text{dom}(r)$ with $r(\beta) = (\alpha, q)$, there is a set S that satisfies $\delta(q, T(\alpha))$; and for each $(i, q') \in S$, there is some j such that $\beta j \in \text{dom}(r)$ and $r(\beta j) = (\alpha i, q')$.

Let $\pi = \pi_1 \pi_2 \dots$ be an infinite path in r ; for each $i \geq 0$, let the state label of the node $\pi_1 \dots \pi_i$ be q_{n_i} where q_{n_0} , the state label of ϵ , is q_I . We say that π satisfies the *parity condition* just if the largest priority that occurs infinitely often in $\Omega(q_{n_0}) \Omega(q_{n_1}) \Omega(q_{n_2}) \dots$ is even. A run-tree r is *accepting* if every infinite path in it satisfies the parity condition. An APT \mathcal{A} accepts a (possibly infinite) ranked tree T if there is an accepting run-tree of \mathcal{A} over T .

Ong [15] has shown that there is a procedure that, given a recursion scheme \mathcal{G} and an APT \mathcal{A} , decides whether \mathcal{A} accepts the value tree of \mathcal{G} .

Theorem 2.1 (Ong). *Let \mathcal{G} be a recursion scheme of order n , and \mathcal{A} be an APT. The problem of deciding whether \mathcal{A} accepts $\llbracket \mathcal{G} \rrbracket$ is n -EXPTIME-complete.*

As usual (following [15]), we restrict our attentions to recursion schemes whose value trees do not contain \perp in the rest of the paper. Given a recursion scheme \mathcal{G} that may generate \perp and an APT \mathcal{A} , one can construct \mathcal{G}' and \mathcal{A}' such that (i) \mathcal{A} accepts $\llbracket \mathcal{G} \rrbracket$ if and only if \mathcal{A}' accepts $\llbracket \mathcal{G}' \rrbracket$, and (ii) \mathcal{G}' does not generate \perp .²

3. TRIVIAL APT AND THE COMPLEXITY OF MODEL CHECKING

APT with a trivial acceptance condition, or *trivial APT* (for short), is an APT that has exactly one priority which is even. Note that trivial APT are equivalent to Aehlig's "trivial automata" [1] (for defining languages of ranked trees).

The first result of this paper is a logical characterization of the class of Σ -labelled ranked trees accepted by trivial APT. Call \mathcal{S} the following fragment of the modal mu-calculus:

$$\phi, \psi ::= \mathbf{t} \mid \mathbf{f} \mid P_f \mid Z \mid \phi \wedge \psi \mid \phi \vee \psi \mid \langle i \rangle \phi \mid \nu Z. \phi$$

where f ranges over symbols in a Σ , and i ranges over $\{1, \dots, \text{arity}(\Sigma)\}$. (We think of \mathcal{S} as the "safety" fragment.) We give a characterization of trivial APT. A proof is given in Appendix A.

Proposition 3.1 (Equi-Expressivity). *The logic \mathcal{S} and trivial APT are equivalent for defining possibly-infinite ranked trees. I.e. for every closed \mathcal{S} -formula, there is a trivial APT that defines the same tree language, and vice versa.*

²Note, however, that the transformation does not preserve the class of trivial APT considered in Section 3.

We show that the model checking problem for recursion schemes is n -EXPTIME complete for trivial APT. The upper-bound of n -EXPTIME follows immediately from Ong's result [15]. To show the lower-bound, we reduce the decision problem of $w \stackrel{?}{\in} \mathcal{L}(\mathcal{A})$, where w is a word and \mathcal{A} is an order- n alternating PDA, to the model checking problem for recursion schemes. n -EXPTIME hardness follows from the reduction, since the problem of $w \stackrel{?}{\in} \mathcal{L}(\mathcal{A})$ is n -EXPTIME hard [6].

Definition 3.2. An *order- n alternating PDA* (order- n APDA, for short) for finite words is a 7-tuple:

$$\mathcal{A} = \langle P, \lambda, p_0, \Gamma, \Sigma, \Delta, F \rangle$$

where P is a set of states, $\lambda : P \rightarrow \{\mathbf{A}, \mathbf{E}\}$ partitions states into universal and existential, p_0 is the initial state, Γ is a stack alphabet, Σ is an input alphabet, $F \subseteq P$ is the set of final states, and $\Delta \subseteq P \times \Gamma \times (\Sigma \cup \{\epsilon\}) \times P \times Op_n$ is a transition relation. A *configuration* of an order- n APDA is of the form (p, s) where s is an order- n stack (an order-1 stack is an ordinary stack, and an order- $(k+1)$ stack is a stack of order- k stacks). The induced transition relation on configurations is defined by the rule:

$$\text{if } (p, top_1(s), \alpha, p', \theta) \in \Delta, \text{ then } (p, s) \longrightarrow_{\alpha} (p', \theta(s))$$

where $\theta \in Op_n$ is an order- n stack operation³ and $top_1(s)$ is the stack top of s .

Let w be a word over Σ . We write w_i (where $0 \leq i < |w|$) for the i -th element of w . A *run tree* of an order- n APDA over a word w is a *finite*, unranked tree satisfying the following.

- (i) The root is labelled by $(p_0, \perp_n, 0)$, where \perp_n is the empty order- n stack.
- (ii) If a node is labelled by (p, s, i) , then one of the following conditions holds, where

$$\begin{aligned} \Xi := & \{(p', \theta(s), i+1) \mid (p, top_1(s), w_i, p', \theta) \in \Delta \wedge i < |w|\} \\ & \cup \{(p', \theta(s), i) \mid (p, top_1(s), \epsilon, p', \theta) \in \Delta\}. \end{aligned}$$

- $p \in F$ and $i = |w|$;
 - $\lambda(p) = \mathbf{A}$ and the set of labels of the child nodes is Ξ ; or
 - $\lambda(p) = \mathbf{E}$ and there is exactly one child node, which is labelled by an element of Ξ .
- (It follows that the leaves of a run tree are labelled by $(p, s, |w|)$ with $p \in F$, or (p, s, i) with $\lambda(p) = \mathbf{A}$ and $\Xi = \emptyset$.)

An order- n APDA \mathcal{A} *accepts* w if there exists a run tree of \mathcal{A} over w .

Engelfriet [6] has shown that the word acceptance problem for order- n APDA is n -EXPTIME complete.

Theorem 3.3 (Engelfriet). *Let \mathcal{A} be an order- n APDA and w a finite word over Σ . The problem of $w \stackrel{?}{\in} \mathcal{L}(\mathcal{A})$ is n -EXPTIME complete.*

³Assume an order- n stack, where $n \geq 2$. An *order-1 push* operation is just the standard operation that pushes a symbol onto the top of the top order-1 stack; the *order-1 pop* operation removes the top symbol from the top order-1 stack. For $2 \leq i \leq n$, the *order- i push* operation duplicates the top order- $(i-1)$ stack of the order- n stack; the *order- i pop* operation removes the top order- $(i-1)$ stack. The set Op_n of order- n stack operations consists of order- i push and order- i pop for each $1 \leq i \leq n$. For a formal definition, see, for example, the FoSSaCS 2002 paper [10] of Knapik et al.

To reduce the word acceptance problem of order- n APDA to the model checking problem for recursion schemes, we use the equivalence [10] between order- n *safe*⁴ recursion schemes and order- n PDA as (deterministic) devices for generating trees.

Definition 3.4. An *order- n tree-generating PDA* is a tuple $\mathcal{A} = \langle \Sigma, \Gamma, Q, \delta, q_0 \rangle$ where Σ is a ranked alphabet, Γ is a stack alphabet, Q is a finite set of states, $q_0 \in Q$ is the initial state, and

$$\delta : Q \times \Gamma \longrightarrow (Q \times Op_n \cup \{(f; q_1, \dots, q_{arity(f)}) \mid f \in \Sigma, q_i \in Q\})$$

is the transition function. A *configuration* is either a pair (q, s) where $q \in Q$ and s is an order- n stack, or a triple of the form $(f; q_1 \dots q_{arity(f)}; s)$ where $f \in \Sigma$ and $q_1 \dots q_{arity(f)} \in Q^*$. Let $\bar{\Sigma}$ be the label-set $\{(f, i) \mid f \in \Sigma, 1 \leq i \leq arity(f)\} \cup \{a \in \Sigma \mid arity(a) = 0\}$. We define the labelled transition relation between configurations induced by δ :

$$\begin{aligned} (q, s) &\xrightarrow{\epsilon} (q', \theta(s)) && \text{if } \delta(q, top_1(s)) = (q', \theta) \\ (q, s) &\xrightarrow{\epsilon} (f; \bar{q}; s) && \text{if } \delta(q, top_1(s)) = (f; \bar{q}) \text{ and } arity(f) \geq 1 \\ (q, s) &\xrightarrow{a} (a; \epsilon; s) && \text{if } \delta(q, top_1(s)) = (a; \epsilon) \text{ and } arity(a) = 0 \\ (f; \bar{q}; s) &\xrightarrow{(f, i)} (q_i, s) && \text{where } 1 \leq i \leq arity(f) \end{aligned}$$

Let w be a finite or infinite word over the alphabet $\bar{\Sigma}$. We say that w is a *trace* of \mathcal{A} just if there is a possibly-infinite sequence of transitions $(q_0, \perp_n) \xrightarrow{\ell_1} \gamma_1 \dots \xrightarrow{\ell_m} \gamma_m \xrightarrow{\ell_{m+1}} \dots$ such that $w = \ell_1 \ell_2 \dots$. We say that \mathcal{A} *generates* a Σ -labelled tree t just in case the branch language⁵ of t coincides with the set of maximal traces of \mathcal{A} .

Theorem 3.5 (Knapik et al. [10]). *There is an effective transformation that, given an order- n tree-generating PDA \mathcal{M} , returns an order- n safe recursion scheme \mathcal{G} that generates the same tree as \mathcal{M} . Moreover, both the running time of the transformation algorithm and the size of \mathcal{G} are polynomial in the size of \mathcal{M} .*

By Theorems 3.3 and 3.5, it suffices to show that, given a word w and an order- n APDA \mathcal{A} , one can construct an order- n tree-generating PDA $\mathcal{M}_{\mathcal{A}, w}$ and a trivial APT \mathcal{B} such that w is accepted by \mathcal{A} if, and only if, the tree generated by $\mathcal{M}_{\mathcal{A}, w}$ is accepted by \mathcal{B} .

Let w be a word over Σ . From w and $\mathcal{A} = \langle P, \lambda, p_0, \Gamma, \Sigma, \Delta, F \rangle$ above, we construct an order- k PDA $\mathcal{M}_{\mathcal{A}, w}$ for generating a $\{A, E, R, T\}$ -labelled tree, which is a kind of run tree of \mathcal{A} over the input word w . The node label **A** (**E**, respectively) means that \mathcal{A} is in a universal (existential, respectively) state; **T** means that \mathcal{A} has accepted the word, and **R** means that \mathcal{A} is stuck (no outgoing transition).

⁴An order- n recursion scheme is *safe* if it satisfies a certain condition called safety [11]. We use the equivalence between safe recursion schemes and higher-order PDA just to prove the lower-bound, so that the knowledge about the safety constraint is not required. See [11, 2] for details of the safety constraint.

⁵The *branch language* of $t : dom(t) \longrightarrow \Sigma$ consists of

- (i) infinite words $(f_1, d_1)(f_2, d_2) \dots$ just if there exists $d_1 d_2 \dots \in \{1, 2, \dots, m\}^\omega$ (where m is the maximum arity of the Σ -symbols) such that $t(d_1 \dots d_i) = f_{i+1}$ for every $i \geq 0$; and
- (ii) finite words $(f_1, d_1) \dots (f_n, d_n) f_{n+1}$ just if there exists $d_1 \dots d_n \in \{1, \dots, m\}^*$ such that $t(d_1 \dots d_i) = f_{i+1}$ for $0 \leq i \leq n$, and the arity of f_{n+1} is 0.

Let N be $\max_{q \in P, a \in \Sigma, \gamma \in \Gamma} |\{(q', a', \theta) \mid (q, \gamma, a', q', \theta) \in \Delta, a' \in \{a, \epsilon\}\}|$. I.e. N is the degree of non-determinacy of \mathcal{A} . We define

$$\mathcal{M}_{\mathcal{A}, w} = \langle \{\mathbf{A} \mapsto N, \mathbf{E} \mapsto N, \mathbf{T} \mapsto 0, \mathbf{R} \mapsto 0\}, \Gamma, Q, \delta, (p_0, 0) \rangle$$

where:

– $Q = (P \times \{0, \dots, |w|\}) \cup \{q_{\top}, q_{\perp}\} \cup (P \times \{0, \dots, |w|\} \times Op_n)$
– $\delta : Q \times \Gamma \longrightarrow (Q \times Op_n \cup \{(g; q_1, \dots, q_k) : g \in \{\mathbf{A}, \mathbf{E}, \mathbf{T}, \mathbf{R}\}, k \geq 0, q_1, \dots, q_k \in Q\})$ is given by:

- (1) $\delta((p, |w|), \gamma) = (\mathbf{T}; \epsilon)$, if $p \in F$
- (2) $\delta((p, i), \gamma) = (\mathbf{A}; (p_1, j_1, \theta_1), \dots, (p_m, j_m, \theta_m), \underbrace{q_{\top}, \dots, q_{\top}}_{N-m})$
if $\lambda(p) = \mathbf{A}$ and $\{(p_1, j_1, \theta_1), \dots, (p_m, j_m, \theta_m)\}$ is:
 $\{(p', i+1, \theta) \mid (p, \gamma, w_i, p', \theta) \in \Delta \wedge i < |w|\} \cup \{(p', i, \theta) \mid (p, \gamma, \epsilon, p', \theta) \in \Delta\}$
- (3) $\delta((p, i), \gamma) = (\mathbf{E}; (p_1, j_1, \theta_1), \dots, (p_m, j_m, \theta_m), \underbrace{q_{\perp}, \dots, q_{\perp}}_{N-m})$
if $\lambda(p) = \mathbf{E}$ and $\{(p_1, j_1, \theta_1), \dots, (p_m, j_m, \theta_m)\}$ is:
 $\{(p', i+1, \theta) \mid (p, \gamma, w_i, p', \theta) \in \Delta \wedge i < |w|\} \cup \{(p', i, \theta) \mid (p, \gamma, \epsilon, p', \theta) \in \Delta\}$
- (4) $\delta((p, i, \theta), \gamma) = ((p, i), \theta)$
- (5) $\delta(q_{\top}, \gamma) = (\mathbf{T}; \epsilon)$
- (6) $\delta(q_{\perp}, \gamma) = (\mathbf{R}; \epsilon)$

Rules (2) and (3) are applied only when rule (1) is inapplicable. $\mathcal{M}_{\mathcal{A}, w}$ simulates \mathcal{A} over the word w , and constructs a tree representing the computation of \mathcal{A} . A state $(p, i) \in P \times \{0, \dots, |w| - 1\}$ simulates \mathcal{A} in state p reading the letter w_i . A state (p, i, θ) simulates an intermediate transition state of \mathcal{A} , where θ is the stack operation to be applied. The states q_{\top} and q_{\perp} are for creating dummy subtrees of nodes labelled with \mathbf{A} or \mathbf{E} , so that the number of children of these nodes adds up to N , the arity of \mathbf{A} and \mathbf{E} . Rule (1) ensures that when \mathcal{A} has read the input word and reached a final state, $\mathcal{M}_{\mathcal{A}, w}$ stops simulating \mathcal{A} and outputs \mathbf{T} . Rule (2) is used to simulate transitions of \mathcal{A} in a universal state, reading the i -th input: $\mathcal{M}_{\mathcal{A}, w}$ constructs a node labelled \mathbf{A} (to record that \mathcal{A} was in a universal state) and spawns threads to simulate all possible transitions of \mathcal{A} . Rule (3) is for simulating \mathcal{A} in an existential state. Note that, if \mathcal{A} gets stuck (i.e. if there is no outgoing transition), all children of the \mathbf{E} -node are labelled \mathbf{R} ; thus failure of the computation can be recognized by the trivial APT given in the following. Rule (4) is just for intermediate transitions. Note that a transition of \mathcal{A} is simulated by $\mathcal{M}_{\mathcal{A}, w}$ in two steps: the first for outputting \mathbf{A} or \mathbf{E} , and the second for changing the stack.

Now we construct a trivial APT \mathcal{B} that accepts the tree generated by $\mathcal{M}_{\mathcal{A}, w}$ if, and only if, w is *not* accepted by \mathcal{A} . The trivial APT \mathcal{B} is given by:

$$\mathcal{B} := \langle \{\mathbf{A}, \mathbf{E}, \mathbf{T}, \mathbf{R}\}, \{q_0\}, \delta, q_0, \{q_0 \mapsto 0\} \rangle$$

where:

$$\delta(q_0, \mathbf{A}) = \bigvee_{i=1}^N (i, q_0) \quad \delta(q_0, \mathbf{E}) = \bigwedge_{i=1}^N (i, q_0) \quad \delta(q_0, \mathbf{T}) = \mathbf{f} \quad \delta(q_0, \mathbf{R}) = \mathbf{t}$$

Intuitively, \mathcal{B} accepts all trees representing a failure computation tree of \mathcal{A} . If the automaton in state q_0 reads \mathbf{T} (which corresponds to an accepting state of \mathcal{A}), it gets stuck. Upon reading \mathbf{A} , the automaton non-deterministically chooses one of the subtrees, and checks

whether the subtree represents a failure computation of \mathcal{A} . On the other hand, upon reading E , the automaton checks that all subtrees represent failure computation trees of \mathcal{A} .

By the above construction, we have:

Theorem 3.6. *Let w be a word, and \mathcal{A} an order- n APDA. Then w is not accepted by \mathcal{A} if, and only if, the tree generated by $\mathcal{M}_{\mathcal{A},w}$ is accepted by \mathcal{B} .*

Corollary 3.7. *The trivial APT acceptance problem for the tree generated by an order- n recursion scheme (i.e. whether the tree generated by a given order- n recursion scheme is accepted by a given trivial APT) is n -EXPTIME hard in the size of the recursion scheme.*

By modifying the encoding, we can also show that the model checking problem is n -EXPTIME-hard in the size of the APT. The idea is to modify $\mathcal{M}_{\mathcal{A},w}$ so that it generates a tree representing computation of \mathcal{A} over not just w but all possible input words, and let a trivial APT check the part of the tree corresponding to the input word w . As a result, the trivial APT depends on the input word w , but the tree-generating PDA does not.

We make the following two assumptions on \mathcal{A} (without loss of generality):

- (i) In each state, if \mathcal{A} can perform an ϵ -transition, then \mathcal{A} cannot perform any input transition i.e. $\{(p', \theta) \mid \exists a \in \Sigma. (p, \gamma, a, p', \theta) \in \Delta\} \neq \emptyset$ implies $\{(p', \theta) \mid (p, \gamma, \epsilon, p', \theta) \in \Delta\} = \emptyset$.
- (ii) There is no transition from a final state i.e. if $p \in F$ then $\{(p', \theta) \mid \exists a \in \Sigma \cup \{\epsilon\}. (p, \gamma, a, p', \theta) \in \Delta\} = \emptyset$.

Given an order- n APDA \mathcal{A} and a word w , we shall construct $\mathcal{M}'_{\mathcal{A}}$ and \mathcal{B}_w , such that w is *not* accepted by \mathcal{A} if, and only if, the tree generated by $\mathcal{M}'_{\mathcal{A}}$ is accepted by \mathcal{B}_w . The difference from the construction of $\mathcal{M}_{\mathcal{A},w}$ and \mathcal{B} above is that $\mathcal{M}'_{\mathcal{A}}$ does not depend on w . The idea is to let $\mathcal{M}'_{\mathcal{A}}$ generate a tree representing the computations of \mathcal{A} over all possible inputs. We then let \mathcal{B}_w traverse the part of the tree corresponding to the computation over w , and check whether the computation is successful.

We define a tree-generating PDA $\mathcal{M}'_{\mathcal{A}} = \langle \Sigma', \Gamma, Q, \delta, q_0 \rangle$ where:

- $\Sigma' = \{\text{Read} \mapsto |\Sigma|, \text{Accept} \mapsto 0, \text{Epsilon} \mapsto 1, \text{A} \mapsto N, \text{E} \mapsto N, \text{T} \mapsto 0, \text{R} \mapsto 0\}$
- $Q = P \cup (P \times (\Sigma \cup \{\epsilon\})) \cup \{q_{\top}, q_{\perp}\} \cup (P \times \text{Op}_k)$
- $q_0 = p_0$
- δ is given by:

$$\begin{aligned}
\delta(p, \gamma) &= (\text{Accept}; \epsilon) \text{ if } p \in F \\
\delta(p, \gamma) &= (\text{Epsilon}; ((p, \epsilon), \text{id})) \text{ if } \{(p', \theta) \mid (p, \gamma, \epsilon, p', \theta) \in \Delta\} \neq \emptyset. \\
\delta(p, \gamma) &= (\text{Read}; ((p, a_1), \text{id}), \dots, ((p, a_n), \text{id})) \\
&\quad \text{if } p \notin F, \{(p', \theta) \mid (p, \gamma, \epsilon, p', \theta) \in \Delta\} = \emptyset \text{ and } \Sigma = \{a_1, \dots, a_n\}. \\
\delta((p, \alpha), \gamma) &= (\text{A}; ((p_1, \theta_1), \dots, (p_m, \theta_m), q_{\top}, \dots, q_{\top})) \\
&\quad \text{if } \lambda(p) = \text{A} \text{ and} \\
&\quad \quad \{(p_1, \theta_1), \dots, (p_m, \theta_m)\} = \{(p', \theta) \mid (p, \gamma, \alpha, p', \theta) \in \Delta\} \\
\delta((p, \alpha), \gamma) &= (\text{E}; ((p_1, \theta_1), \dots, (p_m, \theta_m), q_{\perp}, \dots, q_{\perp})) \\
&\quad \text{if } \lambda(p) = \text{E} \text{ and} \\
&\quad \quad \{(p_1, \theta_1), \dots, (p_m, \theta_m)\} = \{(p', \theta) \mid (p, \gamma, \alpha, p', \theta) \in \Delta\} \\
\delta((p, \theta), \gamma) &= (p, \theta) \\
\delta(q_{\top}, \gamma) &= (\text{T}; \epsilon) \\
\delta(q_{\perp}, \gamma) &= (\text{R}; \epsilon)
\end{aligned}$$

In a final state of \mathcal{A} , $\mathcal{M}'_{\mathcal{A}}$ outputs a node labelled with **Accept**, to indicate that \mathcal{A} has reached a final state, and stops simulating \mathcal{A} (as, by assumption (ii) above, there is no outgoing transition). In a state where \mathcal{A} has ϵ -transitions, $\mathcal{M}'_{\mathcal{A}}$ outputs a node labelled with **Epsilon**, and then simulates all the possible ϵ -transitions of \mathcal{A} . In a state where \mathcal{A} has input transitions, $\mathcal{M}'_{\mathcal{A}}$ outputs a node labelled with **Read** to indicate that \mathcal{A} makes an input transition, and then simulates the input transition for each possible input symbol. Note that by the assumptions (i) and (ii) above, these three transitions are disjoint. The remaining transition rules are analogous to those of $\mathcal{M}_{\mathcal{A},w}$.

Define the trivial APT \mathcal{B}_w by $\mathcal{B}_w = \langle \Sigma', Q', \delta, q_0, \Omega \rangle$ where:

$$\begin{aligned} Q' &= \{q_0, \dots, q_{|w|}\} \\ \delta(q, \mathbf{Epsilon}) &= (1, q) \text{ for every } q \in Q' \\ \delta(q_i, \mathbf{Read}) &= (j, q_{i+1}) \text{ if } 0 \leq i \leq |w| - 1 \text{ and } w_i = a_j \\ \delta(q_{|w|}, \mathbf{Read}) &= \mathbf{t} \\ \delta(q_i, \mathbf{A}) &= (1, q_i) \vee \dots \vee (N, q_i) \\ \delta(q_i, \mathbf{E}) &= (1, q_i) \wedge \dots \wedge (N, q_i) \\ \delta(q_{|w|}, \mathbf{Accept}) &= \mathbf{f} \\ \delta(q_i, \mathbf{Accept}) &= \mathbf{t} \text{ for every } 0 \leq i < |w| \\ \delta(q, \mathbf{T}) &= \mathbf{f} \text{ for every } q \in Q' \\ \delta(q, \mathbf{R}) &= \mathbf{t} \text{ for every } q \in Q' \end{aligned}$$

and Ω is the trivial priority function.

The trivial APT \mathcal{B}_w traverses the tree generated by $\mathcal{M}'_{\mathcal{A}}$ (which represents transitions of \mathcal{A} for all possible inputs), while keeping track of the position of the input head of \mathcal{A} in its state (q_i means that \mathcal{A} is reading the i -th letter of the word w). Upon reading **Read** in state q_i , \mathcal{B}_w proceeds to traverse the branch corresponding to the i -th letter (i.e. w_i). Reading **Accept** in state $q_{|w|}$ means that \mathcal{A} accepts the word w , so that the run of \mathcal{B}_w fails (recall that \mathcal{B}_w accepts the tree just if \mathcal{A} does *not* accept w). Reading **Accept** in state q_i (with $i < |w|$) on the other hand means that \mathcal{A} does not accept w , so that the run of \mathcal{B}_w succeeds. The remaining transition rules are analogous to those of \mathcal{B} .

By the construction above, w is *not* accepted by \mathcal{A} if, and only if, the tree generated by $\mathcal{M}'_{\mathcal{A}}$ is accepted by \mathcal{B}_w . Since only \mathcal{B}_w depends on the input word w , we get:

Theorem 3.8. *The trivial APT acceptance problem of trees generated by order- n recursion schemes is n -EXPTIME-hard in the size of the APT.*

To our knowledge, the lower bound (of the complexity of model-checking recursion schemes) in terms of the size of APT for the entire class of APT is new.

4. DISJUNCTIVE APT AND COMPLEXITY OF MODEL CHECKING

A *disjunctive APT* is an APT whose transition function δ is disjunctive, i.e. δ maps each state to a positive boolean formula θ that contains only disjunctions and no conjunctions, as given by the grammar $\theta ::= \mathbf{t} \mid \mathbf{f} \mid (i, q) \mid \theta \vee \theta$. Disjunctive APT can be used to describe path (or linear-time) properties of trees.

First we give a logical characterization of disjunctive APT as follows. Call \mathcal{D} the following “disjunctive fragment” of the modal mu-calculus:

$$\phi, \psi ::= \mathbf{t} \mid \mathbf{f} \mid P_f \wedge \phi \mid Z \mid \phi \vee \psi \mid \langle i \rangle \phi \mid \nu Z. \phi \mid \mu Z. \phi$$

where f ranges over symbols in Σ , and i over $\{1, \dots, m\}$ where m is the largest arity of the symbols in Σ . A proof of the following proposition is given in Appendix A.

Proposition 4.1 (Equi-Expressivity). *The logic \mathcal{D} and disjunctive APT are equivalent for defining possibly-infinite ranked trees. I.e. for every closed \mathcal{D} -formula, there is a disjunctive APT that defines the same tree language, and vice versa.*

Remark 1. For defining languages of ranked trees, disjunctive APT are a proper subset of the *disjunctive formulas* in the sense of Walukiewicz and Janin [8]. For example, the disjunctive formula $(1 \rightarrow \{t\}) \wedge (2 \rightarrow \{t\})$ is not equivalent to any disjunctive APT.

In the rest of the section, we show that the model checking problem for order- n recursion schemes is $(n - 1)$ -EXPTIME complete for disjunctive APT.

4.1. Upper Bound. Since our proof is based on Kobayashi and Ong's type system for recursion schemes [13] and relies heavily on the machinery and techniques developed therein, we shall just sketch a proof here; a detailed proof will be presented in the journal version of [13]. An alternative proof, also sketched but based on *variable profiles* [15], is given in Appendix B.

Theorem 4.2. *Let \mathcal{G} be an order- n recursion scheme and \mathcal{B} a disjunctive APT. It is decidable in $(n - 1)$ -EXPTIME whether \mathcal{B} accepts the value tree $\llbracket \mathcal{G} \rrbracket$.*

In a recent paper [13], we constructed an intersection type system equivalent to the modal mu-calculus model checking of recursion schemes, in the sense that for every APT, there is a type system such that the tree generated by a recursion scheme is accepted by the APT if, and only if, the recursion scheme is typable in the type system. The model checking problem is thus reduced to a type checking problem. The main idea of the type system is to refine the tree type \circ by the states and priorities of an APT. The type q describes a tree that is accepted by the APT with q as the start state. The type $(\theta_1, m_1) \wedge (\theta_2, m_2) \rightarrow q$, which refines the type $\circ \rightarrow \circ$, describes a tree function that takes an argument which has types θ_1 and θ_2 , and returns a tree of type q .

The type checking algorithm presented in [13] is n -EXPTIME in the combined size of the order- n recursion scheme and the APT (more precisely,⁶ $O(r^{1+\lfloor m/2 \rfloor} \mathbf{exp}_n((a|Q|m)^{1+\epsilon}))$ for $n \geq 2$, where r is the number of rules, a is the largest arity of the symbols in the scheme, m is the largest priority, and $|Q|$ is the number of states). The bottleneck of the algorithm is the number of (atomic) intersection types, where the set $\mathcal{T}(\kappa)$ of atomic types refining a simple type κ is inductively defined by:

$$\begin{aligned} \mathcal{T}(\circ) &:= Q \\ \mathcal{T}(\kappa_1 \rightarrow \kappa_2) &:= \{\bigwedge S \rightarrow \theta \mid \theta \in \mathcal{T}(\kappa_2), S \subseteq \mathcal{T}(\kappa_1) \times P\} \end{aligned}$$

where Q and P are the sets of states and priorities respectively.

According to the syntax of atomic types above, the number of atomic types refining a simple type of order n is n -exponential in general. In the case of disjunctive APT, however, for each type of the form $\circ \rightarrow \dots \rightarrow \circ \rightarrow \circ$, we need to consider only atomic types of the form $\bigwedge S_1 \rightarrow \dots \rightarrow \bigwedge S_k \rightarrow q$, where at most one of the S_i 's is a singleton set and the other S_j 's are empty. Intuitively, this is because a run-tree of a disjunctive APT consists

⁶According to Schewe's recent result [17] on the complexity of parity games, the part $r^{1+\lfloor m/2 \rfloor}$ can be further reduced to roughly $r^{1+m/3}$.

of a single path, so that the run-tree visits only one of the arguments, at most once. In fact, we can show that, if a recursion scheme is typable in the type system for a disjunctive APT, the recursion scheme is typable in a restricted type system in which order-1 types are constrained as described above: this follows from the proof of completeness of the type system [13], along with the property of the accepting run-tree mentioned above. Thus, the number of atomic types is $k \times |Q| \times |P| \times |Q|$ (whereas it is exponential for an arbitrary APT). Therefore, the number of atomic types possibly assigned to a symbol of order n is $(n - 1)$ -exponential. By running the same type checking algorithm as *ibid.* (but with order-1 types constrained as above), order- n recursion schemes can be type-checked (i.e. model-checked) in $(n - 1)$ -EXPTIME.

4.2. Lower Bound. We show the lower bound by a reduction of the emptiness problem of the finite-word language accepted by an order- n deterministic PDA, which is $(n - 1)$ -EXPTIME complete [6].

Let \mathcal{A} be an order- n deterministic PDA, given by $\mathcal{A} = \langle P, p_0, \Gamma, \Sigma, \delta, F \rangle$ where δ is a partial function from $P \times (\Sigma \cup \{\epsilon\}) \times \Gamma$ to $P \times Op_n$. We shall construct an order- n tree-generating PDA $\mathcal{M}_{\mathcal{A}}$, which simulates all possible input and ϵ -transitions of \mathcal{A} , and outputs \mathbf{e} only when \mathcal{A} reaches a final state.

The order- n PDA $\mathcal{M}_{\mathcal{A}}$ is given by:

$$\begin{aligned} \mathcal{M}_{\mathcal{A}} &= \langle \{\mathbf{e} \mapsto 0\} \cup \{\mathbf{br}_m \mapsto m \mid 0 \leq m \leq N\}, \Gamma, P \cup (P \times Op_n), \delta', p_0 \rangle \\ N &= \max_{p \in P, \gamma \in \Gamma} |\{(p', \theta') \mid \exists \alpha \in \Sigma \cup \{\epsilon\}. \delta(p, \alpha, \gamma) = (p', \theta')\}| \\ \delta'(p, \gamma) &= (\mathbf{e}; \epsilon) \text{ if } p \in F \\ \delta'(p, \gamma) &= (\mathbf{br}_m; (p_1, \theta_1), \dots, (p_m, \theta_m)) \\ &\quad \text{if } p \notin F \text{ and } \{(p_1, \theta_1), \dots, (p_m, \theta_m)\} = \{(p', \theta') \mid \exists \alpha \in \Sigma \cup \{\epsilon\}. \delta(p, \alpha, \gamma) = (p', \theta')\} \\ \delta'((p, \theta), \gamma) &= (p, \theta) \end{aligned}$$

A state of $\mathcal{M}_{\mathcal{A}}$ is either a state of \mathcal{A} (i.e. an element of P), or a pair (p, θ) . In state $p \in P$, $\mathcal{M}_{\mathcal{A}}$ constructs a node labeled by \mathbf{br}_m , and spawns subtrees for simulating possible input or ϵ -transitions of \mathcal{A} from state p .

By a result of Knapik et al. [10], we can construct an equi-expressive order- n safe recursion scheme \mathcal{G} . Let \mathcal{G}' be the recursion scheme obtained from \mathcal{G} by (i) replacing each terminal symbol \mathbf{br}_m ($m > 2$) with a non-terminal Br_m of the same arity, and (ii) adding the rule:

$$Br_m x_1 \cdots x_m \rightarrow \mathbf{br}_2 x_1 (\mathbf{br}_2 x_2 (\cdots (\mathbf{br}_2 x_{m-1} x_m))).$$

By the construction, the finite word-language accepted by \mathcal{A} is non-empty if, and only if, the value tree of \mathcal{G}' has a node labelled \mathbf{e} . The latter property can be expressed by the following disjunctive APT \mathcal{B} . (The purpose of transforming \mathcal{G} into \mathcal{G}' was to make the disjunctive APT independent of \mathcal{A} .)

$$\begin{aligned} \mathcal{B} &:= \langle \{q_0\}, \{\mathbf{e}, \mathbf{br}_2\}, \delta, q_0, \{q_0 \mapsto 1\} \rangle \\ &\quad \text{where } \delta(q_0, \mathbf{br}_2) = (1, q_0) \vee (2, q_0) \text{ and } \delta(q_0, \mathbf{e}) = \mathbf{t} \end{aligned}$$

Thus, we have:

Theorem 4.3. *The disjunctive APT acceptance problem for the tree generated by an order- n recursion scheme is $(n - 1)$ -EXPTIME-hard in the size of the recursion scheme.*

The problem is $(n - 1)$ -EXPTIME hard also in the size of the disjunctive APT.

As above, let $\mathcal{A} = \langle P, p_0, \Gamma, \Sigma, \delta, F \rangle$ be an order- n deterministic PDA for words. We may assume that the stack alphabet is $\{\gamma_0, \gamma_1\}$ (as we can encode an arbitrary stack symbol as a sequence of γ_0 and γ_1).

We first define an order- n tree-generating PDA \mathcal{M} by:

$$\begin{aligned} \mathcal{M} &= \langle \{\gamma_0, \gamma_1\}, \{\gamma_0, \gamma_1\}, \{q_0, \theta_1, \dots, \theta_k\}, q_0, \delta_{\mathcal{M}} \rangle \\ \delta_{\mathcal{M}}(q_0, \gamma_i) &= (\gamma_i; \theta_1, \dots, \theta_k) \\ \delta_{\mathcal{M}}(\theta_i, \gamma_j) &= (q_0, \theta_i) \end{aligned}$$

where $\{\theta_1, \dots, \theta_k\}$ is the set of order- n stack operations. The role of \mathcal{M} is to generate a tree simulating all the possible changes of the stack top. Note that \mathcal{M} is independent of \mathcal{A} .

Now let us define a disjunctive APT $\mathcal{D}_{\mathcal{A}} = \langle P, \{\gamma_0, \gamma_1\}, \delta', p_0, \Omega \rangle$ as follows.

$$\begin{aligned} \delta'(p, \gamma_i) &= \begin{cases} \bigvee \{(j, p') \mid \exists \alpha. \delta(p, \gamma_i, \alpha) = (p', \theta_j)\} & \text{if } p \notin F \\ \text{t} & \text{if } p \in F \end{cases} \\ \Omega(p) &= 1 \end{aligned}$$

The idea of the above encoding is to let $\mathcal{D}_{\mathcal{A}}$ simulate transitions of \mathcal{A} , while extracting information about the stack top from the tree generated by \mathcal{M} . Let \mathcal{G} be an order- n recursion scheme that generates the same tree as \mathcal{M} . By the above construction, the language of \mathcal{A} is non-empty if, and only if, $\mathcal{D}_{\mathcal{A}}$ accepts the tree generated by \mathcal{G} . Since the size of \mathcal{G} does not depend on \mathcal{A} , and the size of $\mathcal{D}_{\mathcal{A}}$ is polynomial in the size of \mathcal{A} , we have:

Theorem 4.4. *The disjunctive APT acceptance problem for trees generated by order- n recursion schemes is $(n - 1)$ -EXPTIME hard in the size of the APT.*

4.3. Path Properties. Path properties of Σ -labelled trees are relevant to program verification, as demonstrated in the application to resource usage analysis in Section 5. The *path language* of a Σ -labelled tree t is the image of the map F , which acts on the elements of the branch language of t by “forgetting the argument positions” i.e.

$$F : \begin{cases} (f_1, d_1) (f_2, d_2) \cdots & \mapsto f_1 f_2 \cdots \\ (f_1, d_1) \cdots (f_n, d_n) f_{n+1} & \mapsto f_1 \cdots f_n f_{n+1}^\omega. \end{cases}$$

For example $\{f a^\omega, f f a^\omega, f f b^\omega\}$ is the path language of the term-tree $f a (f a b)$. Let \mathcal{G} be a recursion scheme. We write $\mathcal{W}(\mathcal{G})$ for the *path language* of $\llbracket \mathcal{G} \rrbracket$. Thus elements of $\mathcal{W}(\mathcal{G})$ are infinite words over the alphabet Σ which is now considered unranked (i.e. arities of the symbols are forgotten).

Theorem 4.5. *Let \mathcal{G} be an order- n recursion scheme. The following problems are $(n - 1)$ -EXPTIME complete.*

- (i) $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}) \stackrel{?}{=} \emptyset$, where \mathcal{C} is a non-deterministic parity word automaton.
- (ii) $\mathcal{W}(\mathcal{G}) \stackrel{?}{\subseteq} \mathcal{L}(\mathcal{C})$, where \mathcal{C} is a deterministic parity word automaton.

Furthermore, the problem (i) is $(n - 1)$ -EXPTIME hard not only in the size of \mathcal{G} but also in the size of \mathcal{C} .

Proof. (i) Let $\mathcal{C} = \langle Q, \Sigma, \Delta, q_I, \Omega \rangle$ be a non-deterministic parity word automaton, where $\Delta \subseteq Q \times \Sigma \times Q$ and $\Omega : Q \rightarrow \{0, \dots, p\}$. Let m be the largest arity of the symbols in Σ . (Büchi automata are equivalent to parity automata with two priorities.) We have

$\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}) \neq \emptyset$ if, and only if, $\llbracket \mathcal{G} \rrbracket$ is accepted by the APT $\mathcal{B} = \langle Q, \Sigma, \delta, q_I, \Omega \rangle$ where $\delta : Q \times \Sigma \rightarrow \mathcal{B}^+(\{1, \dots, m\} \times Q)$ is a *disjunctive* transition function

$$\delta : (q, f) \mapsto \bigvee \{(i, p) : 1 \leq i \leq \Sigma(f), (q, f, p) \in \Delta\}.$$

It follows from Theorem 4.2 that the problem $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}) \stackrel{?}{=} \emptyset$ can be decided in $(n-1)$ -EXPTIME.

Let \mathcal{C} be a parity word automaton that accepts $\Sigma^* \mathbf{e}^\omega$, and \mathcal{G}' be the recursion scheme in Section 4.2. Then, $\mathcal{W}(\mathcal{G}') \cap \mathcal{L}(\mathcal{C}) \neq \emptyset$ if, and only if, \mathcal{G}' has a node labelled \mathbf{e} . Thus, the problem $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}) \stackrel{?}{=} \emptyset$ is $(n-1)$ -EXPTIME-hard in the size of \mathcal{G} .

To show the lower bound in the size of \mathcal{C} , we modify the construction of \mathcal{M} and $\mathcal{D}_{\mathcal{A}}$ as follows. Let \mathcal{M}' be the order- n tree-generating PDA given by:

$$\begin{aligned} \mathcal{M}' &:= \langle \{\gamma_0, \gamma_1, \theta_1, \dots, \theta_k\}, \{\gamma_0, \gamma_1\}, \{q_0, q_1, \dots, q_k, \theta_1, \dots, \theta_k\}, q_0, \delta \rangle \\ \delta(q_0, \gamma_i) &= (\gamma_i; q_1, \dots, q_k) \text{ for } 0 \leq i \leq 1 \\ \delta(q_j, \gamma_i) &= (\theta_i; \theta_i) \text{ for } 0 \leq i \leq 1, 1 \leq j \leq k \\ \delta(\theta_j, \gamma_i) &= (q_0, \theta_i) \text{ for } 0 \leq i \leq 1, 1 \leq j \leq k \end{aligned}$$

The difference from \mathcal{M} is that \mathcal{M}' outputs not only stack top symbols but also stack operations (which were coded as branch information in the case of \mathcal{M}). Let $\mathcal{C}_{\mathcal{A}}$ be the non-deterministic parity word automaton given by:

$$\begin{aligned} \mathcal{C}_{\mathcal{A}} &:= \langle P \cup (P \times \{0, 1\}), \{\gamma_0, \gamma_1\}, \delta', p_0, \Omega \rangle \\ \delta'(p, \gamma_i) &= \{(p, i)\} \text{ if } p \notin F \\ \delta'((p, i), \theta_j) &= \{p' \mid \exists \alpha. \delta(p, \gamma_i, \alpha) = (p', \theta_j)\} \\ \delta'(p, \gamma_i) &= \{p\} \text{ if } p \in F \\ \delta'(p, \theta_j) &= \{p\} \\ \Omega(p) &= \begin{cases} 2 & \text{if } p \in F_{\mathcal{A}} \\ 1 & \text{otherwise} \end{cases} \end{aligned}$$

Let \mathcal{G} be a recursion scheme that generates the same tree as \mathcal{M}' . Then, the language of \mathcal{A} is empty if, and only if, $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}_{\mathcal{A}}) = \emptyset$. Since \mathcal{G} does not depend on \mathcal{A} , $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\mathcal{C}) \stackrel{?}{=} \emptyset$ is $(n-1)$ -EXPTIME hard also in the size of \mathcal{C} .

(ii) Let \mathcal{C} be a *deterministic* parity word automaton $\mathcal{C} = \langle Q, \Sigma, \delta_{\mathcal{C}}, q_0, \Omega \rangle$, where $\delta_{\mathcal{C}} : Q \times \Sigma \rightarrow Q$ and $\Omega : Q \rightarrow \{0, \dots, p\}$. Define $\overline{\mathcal{A}} = \langle Q, \Sigma, \delta_{\mathcal{C}}, q_0, \overline{\Omega} \rangle$ where $\overline{\Omega} : q \mapsto (\Omega(q) + 1)$. Note that because of determinacy, $\mathcal{L}(\overline{\mathcal{C}}) = \Sigma^\omega \setminus \mathcal{L}(\mathcal{C})$. Now we have $\mathcal{W}(\mathcal{G}) \subseteq \mathcal{L}(\mathcal{C})$ if, and only if, $\mathcal{W}(\mathcal{G}) \cap \mathcal{L}(\overline{\mathcal{C}}) = \emptyset$. Thus, the problem $\mathcal{W}(\mathcal{G}) \stackrel{?}{\subseteq} \mathcal{L}(\mathcal{C})$ is $(n-1)$ -EXPTIME. Moreover, since the language $\Sigma^* \mathbf{e}^\omega$ is accepted by a deterministic parity word automaton, the problem is also $(n-1)$ -EXPTIME hard (in the size of \mathcal{G}). \square

The decision problems REACHABILITY (i.e. whether $\llbracket \mathcal{G} \rrbracket$ has a node labelled by a given symbol \mathbf{e}) and FINITENESS (i.e. whether $\llbracket \mathcal{G} \rrbracket$ is finite) are instances of Problem (i) of Theorem 4.5; hence they are in $(n-1)$ -EXPTIME (the former is $(n-1)$ -EXPTIME complete, by the proof of Section 4.2).

Consider the problem LTL MODEL-CHECKING:

“Given an LTL-formula ϕ (generated from atomic propositions of the form P_f with $f \in \Sigma$) and an order- n recursion scheme \mathcal{G} , does every path in $\llbracket \mathcal{G} \rrbracket$ satisfy ϕ ? (Precisely, is $\mathcal{W}(\mathcal{G}) \subseteq \llbracket \phi \rrbracket$?)”

As a corollary of Theorem 4.5, we have:

Corollary 4.6. LTL MODEL-CHECKING (i.e. given order- n recursion scheme \mathcal{G} and LTL-formula ϕ , is $\mathcal{W}(\mathcal{G}) \subseteq \llbracket \phi \rrbracket$?) is $(n - 1)$ -EXPTIME complete in the size of \mathcal{G} .

Proof. The upper bound follows from Theorem 4.5(i): note that $\mathcal{W}(\mathcal{G}) \subseteq \llbracket \phi \rrbracket$ is equivalent to $\mathcal{W}(\mathcal{G}) \cap \llbracket \neg\phi \rrbracket = \emptyset$, and because $\llbracket \neg\phi \rrbracket$ is ω -regular, it is recognizable [18] by a parity automaton.

The lower bound follows from the $(n - 1)$ -EXPTIME hardness of REACHABILITY: checking whether a recursion scheme satisfies the formula $G(\neg e)$ is $(n - 1)$ -EXPTIME hard in the size of the recursion scheme. \square

Note however that LTL MODEL-CHECKING is n -EXPTIME in the size of the LTL-formula ϕ , as the size of the corresponding parity word automaton is exponential in ϕ in general [19].

5. APPLICATION TO RESOURCE USAGE VERIFICATION

Now we apply the result of the previous section to show that the resource usage verification problem [7] is $(n - 1)$ -EXPTIME complete. The aim of resource usage verification is to check whether a program accesses each resource according to a given resource specification. For example, consider the following program.

```
let rec g x = if rand() then close(x) else (read(x); g(x)) in
let r = open_in "foo" in g(r)
```

Here, `rand()` returns a non-deterministic boolean. The program first defines a recursive function `g` that takes a file pointer `x` as an argument parameter, closes it after some read operations. The program then opens a read-only file “foo”, and passes it to `g`. For this program, the goal of the verification is to statically check that the file is eventually closed before the program terminates, and after it is closed, it is never read from or written to.

Kobayashi [12] recently showed that the resource usage verification problem is decidable for the simply-typed λ -calculus with recursion, generated from a base type of booleans, and augmented by resource creation/access primitives, by reduction to the model checking problem for recursion schemes. Prior to Kobayashi’s work [12], only sound but incomplete verification methods have been proposed.

Following [12], we consider below a simply-typed, call-by-name functional language with only top-level function definitions and resource usage primitives.⁷ A *program* is a triple (D, S, \mathcal{C}) where D is a set of function definitions, S is a function name (representing the main function), and $\mathcal{C} = (Q_{\mathcal{C}}, \Sigma_{\mathcal{C}}, \delta_{\mathcal{C}}, q_{0,\mathcal{C}}, F_{\mathcal{C}})$ is a deterministic word automaton, which describes how the state of a resource is changed by each access primitive. A function definition is of the form $F \tilde{x} = e$, where e is given by:

$$e ::= \star \mid x \mid F \mid e_1 e_2 \mid \mathbf{If}^* e_1 e_2 \mid \mathbf{New}^q e \mid \mathbf{Acc}_a e_1 e_2$$

The term \star is the unit value. The term $\mathbf{If}^* e_1 e_2$ is a non-deterministic branch between e_1 and e_2 . The term $\mathbf{New}^q e$ creates a fresh resource, and passes it to e (which is a function

⁷Note that programs in call-by-value languages can be transformed into this language by using the standard CPS transformation and λ -lifting.

that takes a resource as an argument). Here, q represents the initial state of a resource; the automaton \mathcal{C} specifies how the resource should be accessed afterwards: see the operational semantics given later. The term $\mathbf{Acc}_a e_1 e_2$ accesses the resource e_1 with the primitive of name $a (\in \Sigma_{\mathcal{C}})$ and then executes e_2 .

Programs must be simply typed; the two base types are **unit** for unit values and **R** for resources. The body of each definition must have type **unit** (in other words, resources cannot be used as return values; this requirement can be enforced by the CPS transformation [16, 5]). The constants \mathbf{If}^* , \mathbf{New}^q , and \mathbf{Acc}_a are given the following types.

$$\mathbf{If}^* : \mathbf{unit} \rightarrow \mathbf{unit} \rightarrow \mathbf{unit}, \mathbf{New}^q : (\mathbf{R} \rightarrow \mathbf{unit}) \rightarrow \mathbf{unit}, \mathbf{Acc}_a : \mathbf{R} \rightarrow \mathbf{unit} \rightarrow \mathbf{unit}$$

Example 5.1. The program given at the beginning of this section can be expressed as (D, S, \mathcal{C}) where

$$\begin{aligned} D &= \{S = \mathbf{New}^{q_1} (G \star), G k x = \mathbf{If}^* (\mathbf{Acc}_c x k) (\mathbf{Acc}_r x (G k x))\} \\ \mathcal{C} &= (\{q_1, q_2\}, \{\mathbf{r}, \mathbf{c}\}, \delta, q_1, \{q_2\}) \\ \delta(q_1, \mathbf{r}) &= q_1 \quad \delta(q_1, \mathbf{c}) = q_2 \end{aligned}$$

Here, G corresponds to the function g in the original program, and the additional parameter k represents a continuation. The automaton \mathcal{C} specifies that the resource should be accessed according to $\mathbf{r}^* \mathbf{c}$.

We introduce the operational semantics to formally define the resource usage verification problem. A run-time state is either an error state **Error** or a pair (ρ, e) where ρ is a finite map from variables to $Q_{\mathcal{C}}$, which represents the state of each resource. The reduction relation $\longrightarrow_{D, \mathcal{C}}$ on run-time states is defined by:

$$\begin{array}{c} \frac{F \tilde{x} = e' \in D}{(\rho, F \tilde{e}) \longrightarrow_{D, \mathcal{C}} (\rho, [\tilde{e}/\tilde{x}]e')} \\ \frac{}{(\rho, \mathbf{If}^* e_1 e_2) \longrightarrow_{D, \mathcal{C}} (\rho, e_1)} \\ \frac{}{(\rho, \mathbf{If}^* e_1 e_2) \longrightarrow_{D, \mathcal{C}} (\rho, e_2)} \\ \frac{x \notin \text{dom}(\rho)}{(\rho, \mathbf{New}^q e) \longrightarrow_{D, \mathcal{C}} (\rho\{x \mapsto q\}, e x)} \\ \frac{\delta_{\mathcal{C}}(q, a) = q'}{(\rho\{x \mapsto q\}, \mathbf{Acc}_a x e) \longrightarrow_{D, \mathcal{C}} (\rho\{x \mapsto q'\}, e)} \\ \frac{\delta_{\mathcal{C}}(q, a) \text{ is undefined}}{(\rho\{x \mapsto q\}, \mathbf{Acc}_a x e) \longrightarrow_{D, \mathcal{C}} \mathbf{Error}} \end{array}$$

Example 5.2. Recall the program in Example 5.1. It can be reduced as follows.

$$\begin{aligned}
(\emptyset, S) &\longrightarrow_{D, \mathcal{C}} (\emptyset, \mathbf{New}^{q_1} (G \star)) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, G \star y) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, \mathbf{If}^* (\mathbf{Acc}_c y \star) (\mathbf{Acc}_r y (G \star y))) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, \mathbf{Acc}_r y (G \star y)) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, G \star y) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, \mathbf{If}^* (\mathbf{Acc}_c y \star) (\mathbf{Acc}_r y (G \star y))) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_1\}, \mathbf{Acc}_c y \star) \\
&\longrightarrow_{D, \mathcal{C}} (\{y \mapsto q_2\}, \star)
\end{aligned}$$

We can now formally define the resource usage verification problem.

Definition 5.3 (resource usage verification problem). A program (D, S, \mathcal{C}) is *resource-safe* if (i) $(\emptyset, S) \not\rightarrow_{D, \mathcal{C}}^* \mathbf{Error}$, and (ii) if $(\emptyset, S) \rightarrow_{D, \mathcal{C}}^* (\rho, \star)$ then $\rho(x) \in F_{\mathcal{C}}$ for every $x \in \text{dom}(\rho)$. The *resource usage verification* is the problem of checking whether a program is resource-safe.

Example 5.4. The program given in Example 5.1 is resource-safe. The program obtained by replacing the body of G (i.e. $\mathbf{If}^* (\mathbf{Acc}_c x k) (\mathbf{Acc}_r x (G k x))$) with $\mathbf{Acc}_r x (G k x)$ is also resource-safe; it does not terminate, so that it satisfies condition (ii) of Definition 5.3 vacuously. The program D' obtained by replacing the definition of G with:

$$G k x = \mathbf{If}^* k (\mathbf{Acc}_r x (G k x))$$

is not resource-safe, as $(\emptyset, S) \rightarrow_{D', \mathcal{C}}^* (\{y \mapsto q_1\}, \star)$ and $q_1 \notin F_{\mathcal{C}}$.

We show below that the resource usage verification is $(n - 1)$ -EXPTIME complete for $n \geq 3$, where n is the largest order of types of terms in the source program. Here, the order of a type is defined by:

$$\text{order}(\mathbf{unit}) = 0 \quad \text{order}(\mathbf{R}) = 1 \quad \text{order}(\kappa_1 \rightarrow \kappa_2) = \max(\text{order}(\kappa_1) + 1, \text{order}(\kappa_2))$$

Note that 3 is the lowest order of a closed program that creates a resource, since \mathbf{New}^q has order 3.

The lower-bound can be shown by reduction of the reachability problem for a recursion scheme to the resource usage verification problem: Given a recursion scheme $\mathcal{G} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$, let (D, S, \mathcal{C}) be the program given by:

$$\begin{aligned}
D &= \{F \tilde{x} = g2p(t) \mid F \tilde{x} \rightarrow t \in \mathcal{R}\} \cup \{\mathbf{Fail} x = \mathbf{Acc}_{\mathbf{fail}} x \star\} \\
g2p(F t_1 \cdots t_m) &= F g2p(t_1) \cdots g2p(t_m) \\
g2p(\mathbf{e}) &= \mathbf{New}^q \mathbf{Fail} \\
g2p(a t_1 \cdots t_m) &= \mathbf{If}^* g2p(t_1) (\cdots (\mathbf{If}^* g2p(t_{m-1}) g2p(t_m))) (a \neq \mathbf{e}) \\
\mathcal{C} &= (\{q\}, \{\mathbf{fail}\}, \emptyset, q, \{q\})
\end{aligned}$$

Then, the value tree of \mathcal{G} contains \mathbf{e} if and only if the program (D, S, \mathcal{C}) is not resource-safe. Since resource primitives occur only in the encoding of \mathbf{e} , the order of the program is the maximum of 3 and the order of the recursion scheme.

To show the upper-bound, we transform a program (D, S, \mathcal{C}) into a recursion scheme $\mathcal{G}_{(D, S, \mathcal{C})}$, which generates a tree representing all possible (resource-wise) access sequences of the program [12], and a disjunctive APT $\mathcal{D}_{(D, S, \mathcal{C})}$, which accepts trees containing an invalid resource access sequence, so that (D, S, \mathcal{C}) is resource-safe if, and only if, $\mathcal{D}_{(D, S, \mathcal{C})}$ rejects the value tree of $\mathcal{G}_{(D, S, \mathcal{C})}$.




Figure 2: The tree generated by the recursion scheme of Example 5.5

The recursion scheme $\mathcal{G}_{(D,S,C)} = (\Sigma, \mathcal{N}, \mathcal{R}, S)$ is given by:

$$\begin{aligned} \Sigma &= \{a \mapsto 1 \mid a \in A\} \cup \{\nu^q \mapsto 2 \mid q \in Q_C\} \cup \{\star \mapsto 0, \mathbf{i} \mapsto 1, \mathbf{k} \mapsto 1, \mathbf{br} \mapsto 2\} \\ \mathcal{N} &= (\text{the set of function symbols in } D) \\ &\quad \cup \{\mathbf{If}^* \mapsto \circ \rightarrow \circ \rightarrow \circ\} \cup \{\mathbf{Acc}_a \mapsto (\circ \rightarrow \circ) \rightarrow \circ \rightarrow \circ \mid a \in A\} \\ &\quad \cup \{\mathbf{New}^q \mapsto ((\circ \rightarrow \circ) \rightarrow \circ) \rightarrow \circ \mid q \in Q_C\} \\ \mathcal{R} &= \{F \tilde{x} \rightarrow e \mid F \tilde{x} = e \in D\} \\ &\quad \cup \{\mathbf{If}^* x y \rightarrow \mathbf{br} x y, \quad \mathbf{Acc}_a x k \rightarrow x(a k), \quad \mathbf{New}^q k \rightarrow \nu^q(k \mathbf{i})(k \mathbf{k})\} \end{aligned}$$

Here, A is the set of the names of access primitives that occur in D .

The preceding encoding is slightly different from the one presented in [12]. The terminal symbol \mathbf{br} represents a non-deterministic choice. In the rule for \mathbf{New}^q , a fresh resource is instantiated to either \mathbf{i} or \mathbf{k} of arity 1. This is a trick used to extract resource-wise access sequences, by tracking or ignoring the new resource in a non-deterministic manner. In the first-branch, the resource is instantiated to \mathbf{i} , so that all the accesses to the resource are kept track of. In the second branch, the resource is instantiated to \mathbf{k} , so that all the accesses to the resource should be ignored. The above transformation preserves types, except that \mathbf{unit} and \mathbf{R} are replaced by \circ and $\circ \rightarrow \circ$ respectively.

Example 5.5. The program in Example 5.1 is transformed into the recursion scheme consisting of the following rules:

$$\begin{aligned} S &\rightarrow \mathbf{New}^{q_1}(G \star) \\ G k x &\rightarrow \mathbf{If}^*(\mathbf{Acc}_c x k)(\mathbf{Acc}_r x (G k x)) \\ \mathbf{If}^* x y &\rightarrow \mathbf{br} x y \\ \mathbf{Acc}_a x k &\rightarrow x(a k) \\ \mathbf{New}^{q_1} k &\rightarrow \nu^{q_1}(k \mathbf{i})(k \mathbf{k}) \end{aligned}$$

Figure 2 shows the value tree of the recursion scheme. The root node represents creation of a new resource (whose initial state is q_1). The nodes labeled by \mathbf{c} or \mathbf{r} express resource accesses. The left and right children are the same, except that each resource access is prefixed by \mathbf{i} in the left child, while it is prefixed by \mathbf{k} in the right child.




Figure 3: The tree generated by the recursion scheme of Example 5.6

Example 5.6. Consider the following program, which creates and accesses two resources:

$$\begin{aligned}
 S &= \mathbf{New}^{q1} F \\
 F x &= \mathbf{New}^{q1} (G \star x) \\
 G k x y &= \mathbf{If}^* (\mathbf{Acc}_c x (\mathbf{Acc}_c y k)) (\mathbf{Acc}_r x (\mathbf{Acc}_r y (G k x y)))
 \end{aligned}$$

It is transformed into the recursion scheme consisting of the following rules:

$$\begin{aligned}
 S &\rightarrow \mathbf{New}^{q1} F \\
 F x &\rightarrow \mathbf{New}^{q1} (G \star x) \\
 G k x y &\rightarrow \mathbf{If}^* (\mathbf{Acc}_c x (\mathbf{Acc}_c y k)) (\mathbf{Acc}_r x (\mathbf{Acc}_r y (G k x y))) \\
 \mathbf{If}^* x y &\rightarrow \mathbf{br} x y \\
 \mathbf{Acc}_a x k &\rightarrow x (a k) \\
 \mathbf{New}^{q1} k &\rightarrow \nu^{q1}(k i) (k k)
 \end{aligned}$$

Figure 3 shows the value tree of the recursion scheme. Of the four subtrees whose roots are labeled by **br**, the leftmost subtree represents accesses to both resources x and y ; in other words, all the accesses to x and y are prefixed by **i**. In the second subtree, only the accesses to x are prefixed by **i**. In the third subtree, only the accesses to y are prefixed by **i**, while in the rightmost subtree, no accesses are prefixed by **i**.

The disjunctive APT $\mathcal{D}_{(D,S,C)} = (\Sigma, Q, \delta, q_I, \Omega)$, which accepts trees having a path corresponding to an invalid access sequence, is given by:

$$\begin{aligned}
Q &= Q_C \cup \{\bar{q} \mid q \in Q_C\} \cup \{q_I\} \\
\delta(q_I, a) &= \begin{cases} (1, q_I) \vee (2, q_I) & \text{if } a = \mathbf{br} \\ (1, q) \vee (2, q_I) & \text{if } a = \nu^q \\ \mathbf{f} & \text{if } a = \star \\ (1, q_I) & \text{otherwise} \end{cases} \\
\delta(q, a) (\text{where } q \in Q_C) &= \begin{cases} (1, q) \vee (2, q) & \text{if } a = \mathbf{br} \\ (1, q) & \text{if } a = \mathbf{i} \\ (1, \bar{q}) & \text{if } a = \mathbf{k} \\ (2, q) & \text{if } a = \nu^q \\ \mathbf{f} & \text{if } a = \star \text{ and } q \in F_C \\ \mathbf{t} & \text{if } a = \star \text{ and } q \notin F_C \\ (1, q') & \text{if } a \in A \text{ and } \delta_C(q, a) = q' \\ \mathbf{t} & \text{if } a \in A \text{ and } \delta_C(q, a) \text{ is undefined} \end{cases} \\
\delta(\bar{q}, a) (\text{where } q \in Q_C) &= (1, q) \\
\Omega(q) &= 1 \text{ for every } q \in Q
\end{aligned}$$

Σ is the same as that of $\mathcal{G}_{(D,S,C)}$.

The APT reads the root of a tree with state q_I , and traverses a tree to find a path corresponding to an invalid resource access sequence. After reading ν^q in state q_I , the APT either (i) chooses the left branch and changes its state to q , the initial state of the new resource, tracking accesses to the resource afterwards; or (ii) chooses the right branch, ignoring accesses to the new resource. In the mode to track resource accesses (i.e., in state $q \in Q$), the APT changes its state according to resource accesses, except: (i) upon reading \mathbf{k} , it skips the next symbol, which represents an access to a resource not being tracked, (ii) upon reading ν^q , it only reads the right branch, ignoring the resource created by this ν^q (as it is already keeping track of another resource), (iii) upon reading $a \in A$ such that $\mathcal{C}(q, a)$ is undefined or reading \star when $q \notin F_C$, it terminates successfully (as an invalid access sequence has been found), and (iv) upon reading \star at state $q \in F_C$, it aborts (as a path being read was actually a valid access sequence). The priority function maps every state to 1, so that no infinite run (that corresponds to an infinite execution sequence of the program without any invalid resource access) is considered an accepting run.

From the construction above, we have:

Theorem 5.7. *(D, S, C) is resource-safe if, and only if, the value tree of $\mathcal{G}_{(D,S,C)}$ is not accepted by $\mathcal{D}_{(D,S,C)}$.*

The proof is similar to the corresponding theorem in [12], hence omitted.⁸

Note that the order of $\mathcal{G}_{(D,S,C)}$ is the same as that of D . Thus, as a corollary of the above theorem and Theorem 4.2, we obtain that the resource usage verification is $(n - 1)$ -EXPTIME.

⁸As mentioned above, the encoding presented in this article is slightly different from the one in [12], but the proofs are similar: they are tedious but rather straightforward.

6. RELATED WORK

Our analysis of the lower bound is based on Engelfriet’s earlier work on the complexity of the iterated pushdown automata word acceptance and emptiness problems, and the results of Knapik et al. on the relationship between higher-order PDA and safe recursion schemes.

The model checking of recursion schemes for the class of trivial APT has been studied by Aehlig [1] (under the name “trivial automata”). He gave a model checking algorithm, but did not discuss its complexity. For the same class, Kobayashi [12] showed that the complexity is linear in the size of recursion schemes, if the types and automata are fixed. For the full modal μ -calculus, Kobayashi and Ong [13] have shown that the complexity is n -EXPTIME in the largest arity of symbols in the recursion scheme, the number of states of the APT, and the largest priority, but polynomial in the number of the rules of the recursion scheme.

Our encoding of the word acceptance problem of an order- n alternating PDA into the model checking problem of an order- n tree-generating PDA (the construction of $\mathcal{M}_{\mathcal{A},w}$ in Section 3) is similar to Cachat and Walukiewicz’s encoding of the word acceptance problem into the reachability game on a higher-order pushdown system [3]. In fact, the tree generated by $\mathcal{M}_{\mathcal{A},w}$ seems to correspond to the unravelling of the game graph of the higher-order pushdown system (where the nodes labelled by E are Player’s positions, and those labelled by A are Opponent’s positions). Thus, n -EXPTIME-hardness of model checking for trivial APT (in the size of the recursion scheme) would follow also from n -EXPTIME hardness of the reachability game on higher-order pushdown systems [3, 4].

7. CONCLUSION

We have considered two subclasses of APT, and shown that the model checking of an order- n recursion scheme is n -EXPTIME complete for trivial APT, and $(n-1)$ -EXPTIME complete for disjunctive APT, both in the size of the recursion scheme and in the size of the APT. As an application, we showed that the resource usage verification problem is $(n-1)$ -EXPTIME complete. The lower bound for the finiteness problem (recall Section 4.3) is left as an open problem.

Acknowledgments. We would like to thank the anonymous reviewers for useful comments. This work was partially supported by Kakenhi 20240001 and EPSRC EP/F036361.

REFERENCES

- [1] K. Aehlig. A finite semantics of simply-typed lambda terms for infinite runs of automata. *Logical Methods in Computer Science*, 3(3), 2007.
- [2] W. Blum and C.-H. L. Ong. Safe lambda calculus. In *Proceedings of the 8th International Conference on Typed Lambda Calculi and Applications (TLCA07)*, pages 39–53. Springer-Verlag, 2007. LNCS 4583.
- [3] T. Cachat and I. Walukiewicz. The complexity of games on higher order pushdown automata. *CoRR*, abs/0705.0262, 2007.
- [4] A. Carayol, M. Hague, A. Meyer, C.-H. L. Ong, and O. Serre. Winning regions of higher-order pushdown games. In *LICS*, pages 193–204, 2008.
- [5] O. Danvy and A. Filinski. Representing control: A study of the cps transformation. *Mathematical Structures in Computer Science*, 2(4):361–391, 1992.
- [6] J. Engelfriet. Iterated stack automata and complexity classes. *Information and Computation*, 95(1):21–75, 1991.

- [7] A. Igarashi and N. Kobayashi. Resource usage analysis. *ACM Transactions on Programming Languages and Systems*, 27(2):264–313, 2005.
- [8] D. Janin and I. Walukiewicz. Automata for the modal mu-calculus and related results. In *Proc. MFCS*, pages 552–562, 1995.
- [9] M. Jurdziński. Small progress measures for solving parity games. In *Proc. STACS*, volume 1770 of *Lecture Notes in Computer Science*, pages 290–301, 2000.
- [10] T. Knapik, D. Niwinski, and P. Urzyczyn. Higher-order pushdown trees are easy. In *FoSSaCS 2002*, volume 2303 of *Lecture Notes in Computer Science*, pages 205–222. Springer-Verlag, 2002.
- [11] T. Knapik, D. Niwiński, and P. Urzyczyn. Higher-order pushdown trees are easy. In *FOSSACS'02*, pages 205–222. Springer, 2002. LNCS Vol. 2303.
- [12] N. Kobayashi. Types and higher-order recursion schemes for verification of higher-order programs. In *Proceedings of ACM SIGPLAN/SIGACT Symposium on Principles of Programming Languages*, pages 416–428, 2009.
- [13] N. Kobayashi and C.-H. L. Ong. A type system equivalent to the modal mu-calculus model checking of higher-order recursion schemes. In *Proceedings of LICS 2009*, pages 179–188. IEEE Computer Society Press, 2009.
- [14] O. Kupferman, M. Y. Vardi, and P. Wolper. An automata-theoretic approach to branching-time model checking. *J. ACM*, 47(2):312–360, 2000.
- [15] C.-H. L. Ong. On model-checking trees generated by higher-order recursion schemes. In *LICS 2006*, pages 81–90. IEEE Computer Society Press, 2006.
- [16] G. D. Plotkin. Call-by-name, call-by-value and the lambda-calculus. *Theor. Comput. Sci.*, 1(2):125–159, 1975.
- [17] S. Schewe. Solving parity games in big steps. In *Proceedings of FSTTCS 2007*, volume 4855 of *Lecture Notes in Computer Science*, pages 449–460. Springer-Verlag, 2007.
- [18] W. Thomas. Languages, automata and logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume 3. Springer-Verlag, 1997.
- [19] M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115:137, 1994.
- [20] I. Walukiewicz. Pushdown processes: games and model-checking. *Information and Computation*, 157:234–263, 2001.

APPENDIX A. CHARACTERIZING TRIVIAL APT AND DISJUNCTIVE APT AS MODAL MU-CALCULUS FRAGMENTS

From Logic to Automata. Consider the following set of modal mu-calculus formulas:

$$\phi, \psi ::= \mathbf{t} \mid \mathbf{f} \mid P_f \mid Z \mid \phi \wedge \psi \mid \phi \vee \psi \mid \langle i \rangle \phi \mid \nu Z. \phi \mid \mu Z. \phi$$

This is a superset of the fragments \mathcal{S} and \mathcal{D} introduced in Sections 3 and 4, respectively.

We can apply the translation of Kupferman et al. [14] to a modal mu-calculus formula to get an equivalent alternating parity tree automaton. We just need to modify the definition of δ ([14, page 339]) by:

$$\begin{aligned} \delta(P_f, g) &= \begin{cases} \mathbf{t} & \text{if } g = f \\ \mathbf{f} & \text{if } g \neq f \end{cases} \\ \delta(\langle i \rangle \phi, g) &= \begin{cases} \text{split}(i, \phi) & \text{if } 1 \leq i \leq \Sigma(g) \\ \mathbf{f} & \text{otherwise} \end{cases} \end{aligned}$$

It is easy to see that the translation maps a \mathcal{S} -formula to a trivial automaton, and a \mathcal{D} -formula to a disjunctive automaton.

From Automata to Logic. Our presentation here follows Walukiewicz [20]. Fix an APT $\mathcal{A} = \langle \Sigma, Q, \delta, q_1, \Omega \rangle$ where $Q = \{q_1, \dots, q_n\}$. Suppose the ordering q_1, \dots, q_n satisfies $\Omega(q_i) \geq \Omega(q_j)$ for every $i < j$. Consider the following n -tuple of modal mu-calculus formulas—call it $\chi_{\mathcal{A}}$ —simultaneously defined by least and greatest fixpoints:

$$\sigma_1 \begin{pmatrix} Z_{11} \\ \vdots \\ Z_{1n} \end{pmatrix} \cdots \sigma_n \begin{pmatrix} Z_{n1} \\ \vdots \\ Z_{nn} \end{pmatrix} \cdot \begin{pmatrix} \chi_1 \\ \vdots \\ \chi_n \end{pmatrix}$$

where $\sigma_i := \mu$ if $\Omega(q_i)$ is odd, and ν otherwise. For each $1 \leq i \leq n$

$$\chi_i := \bigvee_{f \in \Sigma} (P_f \wedge \ulcorner \delta(q_i, f) \urcorner).$$

We define $\ulcorner \delta(q_i, f) \urcorner$ by:

$$\begin{aligned} \ulcorner (d, q_i) \urcorner &:= \langle d \rangle Z_{ii} \\ \ulcorner \mathbf{t} \urcorner &:= \mathbf{t} \\ \ulcorner \mathbf{f} \urcorner &:= \mathbf{f} \\ \ulcorner \varphi_1 \wedge \varphi_2 \urcorner &:= \ulcorner \varphi_1 \urcorner \wedge \ulcorner \varphi_2 \urcorner \\ \ulcorner \varphi_1 \vee \varphi_2 \urcorner &:= \ulcorner \varphi_1 \urcorner \vee \ulcorner \varphi_2 \urcorner \end{aligned}$$

Write $\pi_i(\chi_{\mathcal{A}})$ to be a modal mu-calculus formula (semantically) equivalent to $\chi_{\mathcal{A}}$ projected onto the i -th component (which is well-defined by an application of the Bekič Principle).

Let \mathcal{A} be an APT and t a Σ -labelled ranked tree. Walukiewicz [20] has shown that t is accepted by \mathcal{A} if, and only if, it satisfies $\pi_1(\chi_{\mathcal{A}})$ at the root.

Proposition A.1. (i) If \mathcal{A} is a trivial APT, then $\pi_1(\chi_{\mathcal{A}})$ is a \mathcal{S} -formula.

(ii) If \mathcal{A} is a disjunctive APT, then $\pi_1(\chi_{\mathcal{A}})$ is a \mathcal{D} -formula.

Proof. (i): If \mathcal{A} has only one priority 0, then it follows from the definition that $\chi_{\mathcal{A}}$ is constructed using only ν -fixpoint operator. (ii) Since $\delta(q_i, f)$ is a disjunctive formula, it follows that every conjunction subformula of $\chi_{\mathcal{A}}$ is of the form $P_f \wedge \phi$. \square

APPENDIX B. ALTERNATIVE PROOF OF $(n - 1)$ -EXPTIME UPPER-BOUND FOR
DISJUNCTIVE APT

We sketch an alternative proof of Lemma 4.2, using Ong’s *variable profiles* [15].

In order to appreciate the proof sketched below, some knowledge of the workings of a traversal simulating APT is required. In particular it is necessary to know about *variable profiles* and how they are employed.

Since \mathcal{B} is disjunctive, it has an accepting run-tree on $\llbracket G \rrbracket$ just in case it has an accepting run-tree that does not branch (i.e. each node of the run-tree has at most one child). It follows that \mathcal{B} has an accepting traversal tree if and only if it has an accepting traversal tree that does not branch.

The key observation is that the traversal-simulating APT \mathcal{C} thus need only ‘guess’ one exit point when it reaches a node labelled by a variable of order one, even if its type has arity greater than one. It follows that we can simplify the definition of variable profiles. A profile of a ground-type variable has the shape (x, q, m, \emptyset) where q is a state and m a colour, which is the same as the general case. However a profile of a variable ϕ of a first-order type $\underbrace{o \rightarrow \cdots \rightarrow o}_k \rightarrow o$ now has the shape (ϕ, q, m, c) where c is either empty or a *singleton* set consisting of a profile of a ground-type variable, as opposed to a set of such profiles. The profiles of variables of order two or higher are defined as in the general case. Thus the number of variable profiles of a given order (at least one) is reduced by one level of exponentiation compared to the general case. Now viewing $\mathbf{VP}(A)$ as denoting the set of variable profiles of type A (of order at least one) restricted to containing either empty or singleton interfaces:

$$\sum_{A \text{ order } i \text{ type}} |\mathbf{VP}(A)| = O(\text{exp}_{i-1}(|Gr_G| \times |Q| \times p))$$

where Q is the state space of \mathcal{B} , p is the number of priorities, and Gr_G is the (finite) graph that unravels to the computation tree $\lambda(G)$. The number of nodes in the parity game induced by the traversal-simulating APT \mathcal{C} and the computation tree $\lambda(G)$ will thus also have bound $O(\text{exp}_{n-1}(|Gr_G| \times |Q| \times p))$ and using Jurdziński’s algorithm [9] we have it that the acceptance parity game can be solved in time $O(\text{exp}_{n-1}(|Gr_G| \times |Q| \times p))$. The problem thus lies in $(n - 1)$ -EXPTIME.