

VLDL Satisfiability and Model Checking via Tree Automata^{*}

Alexander Weinert

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
weinert@react.uni-saarland.de

Abstract. We present novel algorithms solving the satisfiability problem and the model checking problem for Visibly Linear Dynamic Logic (VLDL) in asymptotically optimal time via a reduction to the emptiness problem for tree automata with Büchi acceptance. Since VLDL allows for the specification of important properties of recursive systems, this reduction enables the efficient analysis of such systems.

Furthermore, as the problem of tree automata emptiness is well-studied, this reduction enables leveraging the mature algorithms and tools for that problem in order to solve the satisfiability problem and the model checking problem for VLDL.

1 Introduction

Visibly Linear Dynamic Logic (VLDL) [23] is an expressive formalism for specifying properties of recursive systems that allows for an intuitive and modular specification of an important subclass of context-free properties. Although there exist tight bounds on the asymptotical complexity of the satisfiability- and the model checking problem for VLDL properties [23], the upper bounds for both problems are witnessed by algorithms that rely on an intricate reduction of the problems to the emptiness problem for visibly pushdown automata [2], for which tool support is lacking.

We present novel reductions of the problems of VLDL satisfiability and VLDL model checking to the emptiness problem for tree automata [20], yielding algorithms for both problems running in asymptotically optimal time. Moreover, as the emptiness problem for tree automata reduces to the problem of solving two-player games with perfect information [16], which is of great importance in the fields of program verification and program synthesis and enjoys mature tool support, the algorithms yielded by our reductions allow us to leverage this tool support for solving the problems of VLDL satisfiability and VLDL model checking.

VLDL is an extension of Linear Temporal Logic (LTL) [17], the de-facto standard for the specification of properties of non-recursive systems. Although popular, it is lacking in expressivity, as it cannot even express all ω -regular properties. The logic VLDL addresses this shortcoming by guarding the temporal operators of LTL with visibly pushdown automata (VPAs) [2]. A VPA is a pushdown automaton that operates over a predefined partition of an alphabet into calls, returns, and local actions, and has to push (pop) a symbol onto (off) its stack whenever it reads a call (return). Upon processing local actions, the automaton must not touch the stack.

Due to these restrictions, a VLDL formula can be compiled into an equivalent VPA over infinite words of exponential size [23]. As a first step in this construction, the VLDL formula is translated into a 1-AJA [5], an automaton without stack that is able to jump from a call to its matching return. This automaton can then be transformed into a VPA of exponential size [5]. Since each visibly pushdown automaton is a classical pushdown automaton, the emptiness problem for VPAs is decidable in polynomial time [2]. The translation from 1-AJAs to VPAs, however, is quite involved, as it works for a far more complex model than is needed for the translation of VLDL formulas into 1-AJAs, thus hampering efforts towards an implementation of the translation from VLDL to VPAs. This effort is further encumbered by the scant availability of emptiness checkers and model checkers for pushdown systems.

In this work, we introduce novel algorithms solving both the emptiness problem and the model checking problem for VLDL formulas in asymptotically optimal time using a translation of VLDL formulas to nondeterministic tree automata with Büchi acceptance. The technical core of this translation is formed by an encoding of words over visibly pushdown alphabets into trees that is adapted from the encoding of such words given by Alur and Madhusudan [2], as well as by a translation of the 1-AJAs constructed

^{*} Supported by the project “TriCS” (ZI 1516/1-1) of the German Research Foundation (DFG).

from VLDL formulas into tree automata using an adaptation of the breakpoint-construction by Miyano and Hayashi [15] in order to remove alternation and obtain a nondeterministic automaton. Satisfiability of a VLDL formula is then checked by checking the resulting tree automaton for emptiness. For model checking a visibly pushdown system against a VLDL specification, we translate the negation of the specification as well as the visibly pushdown system into tree automata, which we intersect and check for emptiness.

Thus, we reduce both the satisfiability- and the model checking problem for VLDL to the emptiness problem for nondeterministic tree automata with Büchi acceptance. Hence, we reduce the complex formalism of VLDL to the simple model of nondeterministic tree automata. Moreover, since the problem of tree automata emptiness reduces to that of solving Büchi games, which is solvable efficiently [8, 7] and enjoys mature tool support [10, 11], our novel reductions enable an efficient implementation of satisfiability checkers and model checkers for VLDL.

Related Work There exist a number of logics other than VLDL that capture the class of visibly pushdown languages, most prominently VLTL [6], a fixed-point logic [5] and monadic second order logic augmented with a binary matching predicate (MSO_μ) [2]. We focus here on the logic VLDL, as it most naturally extends the concepts used by LTL [17], the de-facto standard for the specification of non-recursive properties.

Moreover, there exist tools for model checking recursive problems, e.g., BEBOP [3, 4] and MOPED [18, 19]. These tools are, however, no longer under active development, and have, to the best of our knowledge, not found widespread adoption. In combination with the intricate translation of alternating automata into VPAs, this motivates the development of the novel translation of VLDL formulas into tree automata presented in this work.

A number of problems have been reduced to the emptiness problem for tree automata, as they are a natural model for capturing the branching-time behavior of systems [16]. Moreover, the theory of tree automata is well-studied, with its most famous result being equivalence of tree automata and monadic second order logic of two successors [21, 24]. Finally, the emptiness problem for tree automata with Büchi acceptance reduces to the problem of solving two-player Büchi games with perfect information [9]. Such games can be solved efficiently [7] and, since Büchi games are a special case of the ubiquitous parity games, there exists mature tool support for solving them [10, 11].

Our Contributions Firstly, in Section 3 we adapt the tree-encoding of words over visibly pushdown alphabets first introduced by Alur and Madhusudan [2] and show that the resulting trees are recognizable by a tree automaton with Büchi acceptance condition in Theorem 1.

Secondly, in Section 4, we show how to construct tree automata recognizing the encodings of all words satisfying a given VLDL formula in Theorem 2. Moreover, we show that the resulting automaton is of exponential size measured in the size of the original formula and we show that this translation yields an asymptotically optimal algorithm for satisfiability checking of VLDL formulas.

Finally, in Section 5 we provide a translation of visibly pushdown systems into tree automata recognizing the encodings of all traces of the system. When combined with the previously presented translation of VLDL formulas into tree automatas, we obtain an asymptotically optimal algorithm for model checking visibly pushdown systems against VLDL specifications. This result is given in Theorem 3.

2 Preliminaries

In this section we introduce the basic notions used in the remainder of this work, namely (nondeterministic) visibly pushdown automata and related concepts.

2.1 Visibly Pushdown Languages

A pushdown alphabet $\tilde{\Sigma} = (\Sigma_c, \Sigma_r, \Sigma_l)$ is a finite set Σ that is partitioned into calls Σ_c , returns Σ_r and local actions Σ_l . We write $w = w_0 \cdots w_n$ and $\alpha = \alpha_0 \alpha_1 \alpha_2 \cdots$ for finite and infinite words, respectively, and define the stack height reached by any automaton after reading w by $sh(w)$ inductively as $sh(\varepsilon) = 0$, $sh(wc) = sh(w) + 1$ for $c \in \Sigma_c$, $sh(wr) = \max\{0, sh(w) - 1\}$ for $r \in \Sigma_r$, and $sh(wl) = sh(w)$ for $l \in \Sigma_l$. Let α be a finite or infinite word. We say that a call $\alpha_k \in \Sigma_c$ at some position k of α is matched if there

exists a $k' > k$ such that $\alpha_{k'} \in \Sigma_r$ and $sh(\alpha_0 \cdots \alpha_{k-1}) = sh(\alpha_0 \cdots \alpha_{k'})$ and call the return at the smallest such position k' the matching return of c . Otherwise we call c an unmatched call. If α_k is a matched call with $\alpha_{k'}$ as its matching return, we call the infix $\alpha_{k+1} \cdots \alpha_{k'-1}$ of α the nested infix of position k . A word is well-matched if it does not contain a return that is not a matching return.

A visibly pushdown system (VPS) $\mathcal{S} = (Q, \tilde{\Sigma}, \Gamma, \Delta, q_I)$ consists of a finite set Q of states, a pushdown alphabet $\tilde{\Sigma}$, a stack alphabet Γ , which contains a stack-bottom marker \perp , a transition relation $\Delta \subseteq (Q \times \Sigma_c \times Q \times (\Gamma \setminus \{\perp\})) \cup (Q \times \Sigma_r \times \Gamma \times Q) \cup (Q \times \Sigma_l \times Q)$, and an initial state $q_I \in Q$. A configuration (q, γ) of \mathcal{S} is a pair of a state $q \in Q$ and a stack content $\gamma \in \Gamma_c = (\Gamma \setminus \{\perp\})^* \cdot \perp$. The VPS \mathcal{S} induces the configuration graph $G_{\mathcal{S}} = (Q \times \Gamma_c, E)$ with $E \subseteq ((Q \times \Gamma_c) \times \Sigma \times (Q \times \Gamma_c))$ and $((q, \gamma), a, (q', \gamma')) \in E$ if and only if either

- $a \in \Sigma_c$, $(q, a, q', A) \in \Delta$, and $A\gamma = \gamma'$,
- $a \in \Sigma_r$, $(q, a, \perp, q') \in \Delta$, and $\gamma = \gamma' = \perp$,
- $a \in \Sigma_r$, $(q, a, A, q') \in \Delta$, $A \neq \perp$, and $\gamma = A\gamma'$, or
- $a \in \Sigma_l$, $(q, a, q') \in \Delta$, and $\gamma = \gamma'$.

For an edge $e = ((q, \gamma), a, (q', \gamma'))$, we call a the label of e . A run $\pi = (q_0, \gamma_0) \cdots (q_n, \gamma_n)$ of \mathcal{S} on $w = w_0 \cdots w_{n-1} \in \Sigma^*$ is a sequence of configurations where $q_0 = q_I$ and where $((q_i, \gamma_i), w_i, (q_{i+1}, \gamma_{i+1})) \in E$ in $G_{\mathcal{S}}$ for all $i \in [0; n-1]$. Infinite runs of \mathcal{S} on infinite words are defined similarly. We define $traces(\mathcal{S})$ as the set of all infinite words α for which there exists a run of \mathcal{S} on α . Moreover, we define $|\mathcal{S}| = |Q|$.

2.2 Visibly Linear Dynamic Logic

Let P be a finite set of atomic propositions and let $\tilde{\Sigma} = (\Sigma_c, \Sigma_r, \Sigma_l)$ be a partition of $\Sigma = 2^P$. The syntax of VLDL [23] is defined by the grammar $\varphi := p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle \mathfrak{A} \rangle \varphi \mid [\mathfrak{A}] \varphi$, where $p \in P$ and \mathfrak{A} ranges over testing visibly pushdown automata (TVPA) over the fixed alphabet $\tilde{\Sigma}$. A TVPA $\mathfrak{A} = (Q, \tilde{\Sigma}, \Gamma, \Delta, q_I, Q_F, t)$ consists of a VPS $\mathcal{S} = (Q, \tilde{\Sigma}, \Gamma, \Delta, q_I)$, a set of final states $Q_F \subseteq Q$, and a function t mapping states to VLDL formulas over $\tilde{\Sigma}$. We define $|\varphi|$ as the sum of $|\text{cl}(\varphi)|$ and the sum of the numbers of states of the automata contained in φ , where $\text{cl}(\varphi)$ is the set of all subformulas of φ , including those contained as tests in automata and their subformulas. We require this relation subformula-relation to be noncircular. A run of \mathfrak{A} on a finite word w is a run of the underlying VPS \mathcal{S} on w . Such a run is accepting if its final state is in Q_F .

Let φ be a VLDL formula, let $\alpha = \alpha_0 \alpha_1 \alpha_2 \cdots \in \Sigma^\omega$ and let $k \in \mathbb{N}$ be a position in α . We define the semantics of φ in the straightforward way for atomic propositions and Boolean connectives. Furthermore, we define

- $(\alpha, k) \models \langle \mathfrak{A} \rangle \varphi$ if there exists $k' \geq k$ s.t. $(k, k') \in \mathcal{R}_{\mathfrak{A}}(\alpha)$ and $(\alpha, k') \models \varphi$,
- $(\alpha, k) \models [\mathfrak{A}] \varphi$ if for all $k' \geq k$, $(k, k') \in \mathcal{R}_{\mathfrak{A}}(\alpha)$ implies $(\alpha, k') \models \varphi$,

with

$$\mathcal{R}_{\mathfrak{A}}(\alpha) := \{(k, k') \in \mathbb{N} \times \mathbb{N} \mid \exists \text{ acc. run } (q_0, \sigma_0) \cdots (q_{k'-k}, \sigma_{k'-k}) \\ \text{of } \mathfrak{A} \text{ on } \alpha_k \cdots \alpha_{k'-1} \text{ and } \forall m \in [0; k' - k]. (\alpha, k' + m) \models t(q_m)\} .$$

We write $\alpha \models \varphi$ as a shorthand for $(\alpha, 0) \models \varphi$ and say that α is a model of φ in this case. The language of φ is defined as $L(\varphi) = \{\alpha \in \Sigma^\omega \mid \alpha \models \varphi\}$. If $L(\varphi) \neq \emptyset$, we say that φ is satisfiable.

2.3 Tree Automata

Let $\mathbb{B} = \{0, 1\}$ and let Σ be an alphabet. A Σ -tree t is a mapping $t: \mathbb{B}^* \rightarrow \Sigma$. We call a finite word $b \in \mathbb{B}^*$ a node and an infinite word $\beta \in \mathbb{B}^\omega$ a branch. Given a node b , we call the nodes $b0$ and $b1$ the left- and right-hand children of b . Analogously, we call the trees rooted at the left- and right-hand children of b the left- and right-hand subtrees of b , respectively. Moreover, b is the parent of both $b0$ and $b1$. We call the node at address ε the root of t . We say that a branch β contains a node b if b is a prefix of β . Similarly, as each node b is associated with the unique path from the root of the tree to b , we say that a node b' is on the path to b if b' is a prefix of b . If $t(b) = a$, we say that b is labeled with a . Moreover, given a tree t and a node b , we define the sub-tree $t|_b$ of t rooted at b by $t|_b(b') = t(bb')$.

A tree automaton (with Büchi acceptance) $\mathfrak{T} = (Q, \Sigma, \Delta, q_I, Q_F)$ consists of a finite set of states Q , an alphabet Σ , a transition relation $\Delta \subseteq Q \times \Sigma \times Q \times Q$, an initial state $q_I \in Q$, and a set of accepting

states $Q_F \subseteq Q$. A run r of \mathfrak{T} on a Σ -tree t is a Q -tree with $r(\varepsilon) = q_I$ and $(r(b), t(b), r(b0), r(b1)) \in \Delta$ for all $b \in \mathbb{B}^*$. A branch of r is accepting if it contains infinitely many nodes b such that $r(b) \in F$. A run is accepting if all of its branches are accepting, while an automaton \mathfrak{T} accepts a tree t if there exists an accepting run of \mathfrak{T} on t . The language $L(\mathfrak{T})$ of \mathfrak{T} is defined as the set of all trees accepted by \mathfrak{T} . A set of trees is regular if there exists a tree automaton recognizing it. We define $|\mathfrak{T}| = |Q|$. Tree automata are closed under intersection via an adaptation of the construction for the intersection of automata on words. Hence, for tree automata $\mathfrak{T}_1, \mathfrak{T}_2$ there exists a tree automaton \mathfrak{T} with $|\mathfrak{T}| \in \mathcal{O}(|\mathfrak{T}_1| |\mathfrak{T}_2|)$ such that $\mathcal{L}(\mathfrak{T}) = \mathcal{L}(\mathfrak{T}_1) \cap \mathcal{L}(\mathfrak{T}_2)$.

3 Stack Trees

Alur and Madhusudan showed how to encode words over some visibly pushdown alphabet as a tree by “folding away” the nested infixes of calls into subtrees, thus moving a call and its matching return next to each other in the resulting tree [2]. In this section, we slightly adapt their encoding in order to simplify our construction of tree automata later on in Section 4. In that section, we construct for each VLDL formula φ a tree automaton that accepts precisely the encodings of words satisfying φ .

For the remainder of this work, we fix some pushdown alphabet $\tilde{\Sigma} = (\Sigma_c, \Sigma_r, \Sigma_l)$ as a partition of some alphabet Σ . Let $\alpha \in \Sigma^\omega$ be an infinite word and define $\Sigma_\perp = \Sigma \cup \{\perp\}$, where \perp is some fresh symbol. Intuitively, every node in the resulting tree denotes either one position of α , or it is labeled with the special symbol \perp . For a given word $\alpha \in \Sigma^\omega$, we define the function st mapping finite and infinite words over Σ to infinite Σ_\perp -trees in Figure 1. At every matched call, we encode its matched infix and the suffix starting at and including its matched return in the right- and left-hand subtrees, respectively. At an unmatched call, we encode the suffix of the word starting at the symbol succeeding the unmatched call in the right-hand subtree. If the current letter is not a call, we encode the suffix starting at the current letter’s successor in the left-hand subtree. All vertices not encoding a symbol of α are labeled with \perp .

$$\begin{array}{ll}
 \text{st}(cwr\alpha) = \begin{array}{c} c \\ \swarrow \quad \searrow \\ \text{st}(r\alpha) \quad \text{st}(w) \end{array} & \begin{array}{l} \text{if } c \in \Sigma_c \text{ and } r \\ \text{is matching} \\ \text{return of } c \end{array} & \text{st}(c\alpha) = \begin{array}{c} c \\ \swarrow \quad \searrow \\ \text{st}(\varepsilon) \quad \text{st}(\alpha) \end{array} & \begin{array}{l} \text{if } c \in \Sigma_c \text{ and } c \\ \text{is unmatched} \end{array} \\
 \text{st}(x\alpha) = \begin{array}{c} x \\ \swarrow \quad \searrow \\ \text{st}(\alpha) \quad \text{st}(\varepsilon) \end{array} & \begin{array}{l} \text{if } x \in \Sigma_l \cup \Sigma_r \end{array} & \text{st}(\varepsilon) = \begin{array}{c} \perp \\ \swarrow \quad \searrow \\ \text{st}(\varepsilon) \quad \text{st}(\varepsilon) \end{array} &
 \end{array}$$

Fig. 1. Definition of $\text{st}: \Sigma^\omega \cup \Sigma^* \rightarrow T_{\Sigma_\perp}$.

A tree t is a stack tree if $t = \text{st}(\alpha)$ for some $\alpha \in \Sigma^\omega$. We define the set of all stack trees over Σ as $\text{st}(\Sigma^\omega) = \{t(\alpha) \mid \alpha \in \Sigma^\omega\}$.

Theorem 1. *The set $\text{st}(\Sigma^\omega)$ is regular.*

Proof. We first introduce some notation. Let t be a Σ_\perp -tree. We say that a node b is a matched call if $t(b) \in \Sigma_c$ and $t(b0) \in \Sigma_r$. Similarly, $b0$ is a matched return if we have $t(b) \in \Sigma_c$. If all calls and returns in t are matched, we say that t is well-matched. Furthermore, we call a branch β finite in t if it eventually only contains \perp -labeled vertices. Otherwise, we call β infinite in t . Finally, we call a tree finite if all of its branches are finite.

We claim that a Σ_\perp -tree t is a stack tree if and only if $t(\varepsilon) \neq \perp$, if there exists a single branch that is infinite in t , and if the following properties hold true for all $b \in \mathbb{B}^*$:

1. If $t(b) = \perp$, then $t(b0) = t(b1) = \perp$,
2. if $t(b) \in \Sigma_c$ and b is matched, then $t|_{b1}$ is finite and well-matched,
3. if $t(b) \in \Sigma_c$ and b is unmatched, then $t(b0) = \perp$ and $t|_{b1}$ contains no unmatched returns, and

4. if $t(b) \in \Sigma_l \cup \Sigma_r$, then $t(b1) = \perp$.

Note that each of these properties can be checked by a tree automaton. As tree automata are closed under intersection, there also exists a single tree automaton that checks all of the above properties.

It remains to show that the conditions above indeed characterize stack trees, i.e., that a Σ_\perp -tree is a stack tree if and only if it satisfies the conditions above. First note that for all $\alpha \in \Sigma^\omega$, the tree $\text{st}(\alpha)$ clearly satisfies the above conditions. Hence, we now show that for each tree t satisfying these conditions there exists a word $\alpha \in \Sigma^\omega$ such that $t = \text{st}(\alpha)$. We construct such a word via a preorder traversal of t that visits right-hand children before left-hand ones.

We first show how to encode finite-trees, as such trees encode nested infixes of matched calls. Let t be a finite Σ_\perp -tree satisfying conditions 1 through 4 and let $\text{HEIGHT}(t)$ be the minimal k such that for all nodes b with $|b| \geq k$ we have $t(b) = \perp$ and. We construct a word $w \in \Sigma^*$ such that $t = \text{st}(w)$ by induction over $\text{HEIGHT}(t)$. If $\text{HEIGHT}(t) = 0$, then $t(b) = \perp$ for all $b \in \mathbb{B}^*$ due to Condition 1 and thus, $t = \text{st}(\varepsilon)$. If, however, $\text{HEIGHT}(t) > 0$, then first note that $t(\varepsilon) \notin \Sigma_r$, since ε would be an unmatched return in that case. Thus, first assume $t(\varepsilon) = l \in \Sigma_l$. Then $t(1) = \perp$ and $t' = t|_0$ is a well-matched Σ_\perp -tree with $\text{HEIGHT}(t') < \text{HEIGHT}(t)$. Hence, there exists a word $w' \in \Sigma^*$ such that $\text{st}(w') = t'$. Thus, we pick $w = lw'$ and obtain $t = \text{st}(w)$. Now assume $t(\varepsilon) = c \in \Sigma_c$. As every call in t is matched, we obtain $t(0) = r \in \Sigma_r$ and thus, $t|_{01}(b) = \perp$ for all $b \in \mathbb{B}^*$, while there exist words w_1 and w_{00} such that $\text{st}(w_1) = t|_1$ and $\text{st}(w_{00}) = t|_{00}$ due to the induction hypothesis and due to the second condition given above. Hence, we pick $w = cw_1rw_{00}$ and obtain $t = \text{st}(w)$.

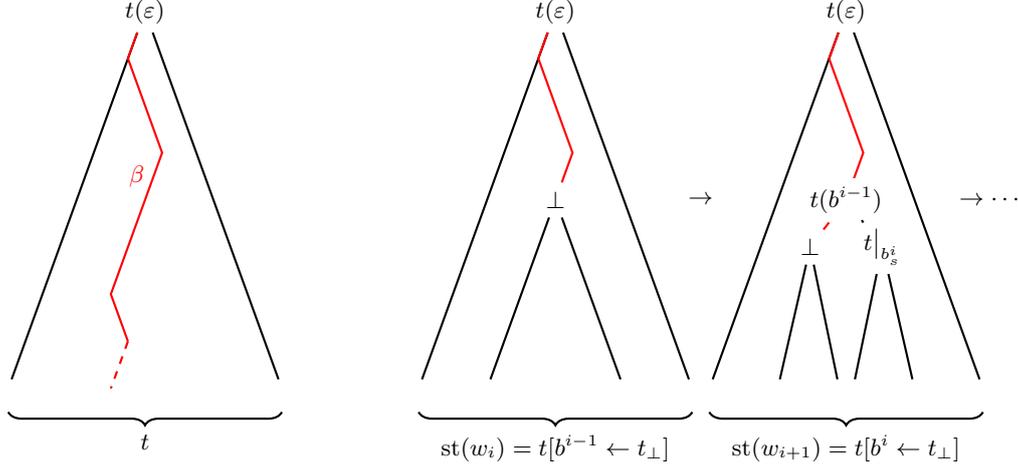


Fig. 2. Construction of the $\text{st}(w_i)$. We have $\beta = b_0b_1b_2 \dots$ and use the shorthands $b^i = b_0 \dots b_i$, and $b_s^i = b_0 \dots b_{i-1}(1 - b_i)$.

Now let t be a tree that satisfies conditions 1 through 4 with $t(\varepsilon) \neq \perp$ and let $\beta = b_0b_1b_2 \dots$ be the single infinite branch of t . As a shorthand, let $b^i = b_0 \dots b_i$ and let t_a be the unique Σ_\perp -tree with $t_a(\varepsilon) = a$ and $t_a(b) = \perp$ for all $b \in \mathbb{B}^+$. Moreover, if t and t' are Σ_\perp -trees and $b \in \mathbb{B}^*$, we define $t[b \leftarrow t']$ such that $t[b \leftarrow t'](b') = t'(b'')$ if $b' = bb''$ for some $b'' \in \mathbb{B}^*$, and $t[b \leftarrow t'] = t(b)$ otherwise. Intuitively, we replace the subtree of t anchored at b by the tree t' . We construct a series of words $w_0, w_1, w_2, \dots \in \Sigma^*$ such that for each w_i ,

1. w_i is a strict prefix of w_{i+1}
2. $\text{st}(w_i) = t[b^{i-1} \leftarrow t_\perp]$, and
3. $\text{st}(w_i \cdot a) = t[b^{i-1} \leftarrow t_a]$, where $a = t(b^{i-1})$.

We illustrate this construction in Figure 2. Due to the first condition, the limit of the w_i for $i \rightarrow \infty$ is an ω -word α , which, due to the second condition, satisfies $\text{st}(\alpha) = t$. The final condition allows us to construct the w_i inductively by collecting the labels of the nodes along the infinite path β of t : Upon encountering a matched return r we are able to append r to the w_i constructed so far and ensure that the resulting w_{i+1} indeed satisfies the second condition.

Formally, we first pick $w_0 = \varepsilon$, which obviously satisfies the above requirements. Now let $i \in \mathbb{N}$ such that w_i is defined and satisfies the above requirements. In order to construct w_{i+1} , let $a = t(b^{i-1})$.

If $a \in \Sigma_c$ and b^{i-1} is matched, then the subtree rooted at b^{i-1} is finite and well-matched due to Condition 2, hence there exists a word $w \in \Sigma^*$ such that $\text{st}(w) = t|_{b^{i-1}}$ as shown above. Thus, it is easy to verify that $w_{i+1} = w_i a w$ satisfies the requirements above. In particular the third requirement is satisfied due to $b_i = 0$, $t(b^i) \in \Sigma_r$, and due to the fact that w is well-matched, i.e., it does not contain unmatched calls.

If $a \in \Sigma_l \cup \Sigma_r$, however, then $w_{i+1} = w_i a$ clearly satisfies the conditions above. In particular, the third condition is satisfied due to b^{i-1} being part of the unique infinite branch β . Hence, $t(b^i) \in \Sigma_r$ can only hold true if there exists no unmatched call on the path to b^i . Thus, the third condition is indeed satisfied. \square

From the proof of Theorem 1 we furthermore obtain that for each $\alpha \in \Sigma^\omega$, there exists exactly one branch $\beta = b_0 b_1 b_2 \dots$ such that $\text{st}(\alpha)(b_0 \dots b_{i-1}) \neq \perp$ for each $i \in \mathbb{N}$. We call β the cardinal branch of $\text{st}(\alpha)$ and we call the positions of the symbols encoded along β the cardinal positions of α .

We give an example of $\text{st}(\alpha)$ for $\alpha = lclrcl\dots$ over the alphabet $\tilde{\Sigma} = (\Sigma_c, \Sigma_r, \Sigma_l) = (\{c\}, \{r\}, \{l\})$ in Figure 3. The positions 0, 1, 3, 4 are cardinal positions, if we assume the second c in α to be unmatched. Recall that we defined $sh(w)$ to be the stack height reached by any visibly pushdown automaton after processing $w \in \Sigma^*$. In general, Löding et al. [14] defined the steps of a word $\alpha = \alpha_0 \alpha_1 \alpha_2 \dots$ as those positions of α that reach a lower bound on the stack height reached during processing the remainder of the word, i.e., $\text{steps}(\alpha) = \{k \mid \forall k' \geq k. sh(\alpha_0 \dots \alpha_k) \leq sh(\alpha_0 \dots \alpha_{k'})\}$. A position k is a cardinal position of α if and only if it is either a steps, or if α_k is the matching return of some call occurring at a step.

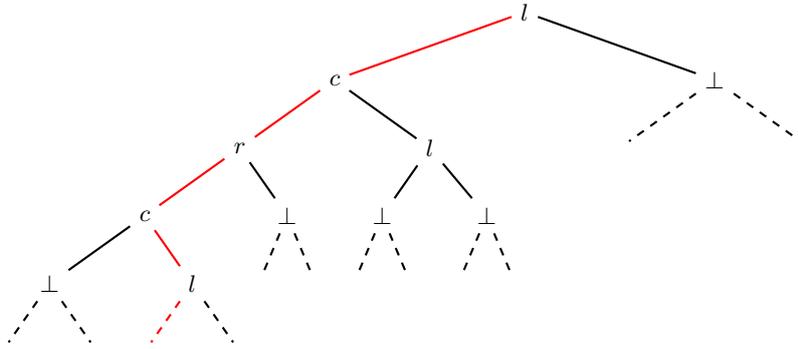


Fig. 3. Encoding of $\alpha = lclrcl\dots$, where the second occurrence of c is an unmatched call. The cardinal branch of $\text{st}(\alpha)$ is marked in red.

4 Reducing VLDL Satisfiability to Tree Automata Emptiness

In this section we reduce the problem of VLDL satisfiability to the emptiness problem for tree automata. The former problem is formulated as follows: “Given some VLDL formula φ , is φ satisfiable?” We formalize the reduction of this problem to the emptiness problem for tree automata as follows:

Theorem 2. *For every VLDL formula φ there exists an effectively constructible tree automaton \mathfrak{T} such that $\mathcal{L}(\mathfrak{T}) = \text{st}(\mathcal{L}(\varphi))$ with $|\mathfrak{T}| \in \mathcal{O}(2^{|\varphi|})$.*

Due to Theorem 2, we obtain an algorithm that checks VLDL formulas for satisfiability by first transforming a given formula φ into the tree automata \mathfrak{T} recognizing $\text{st}(\mathcal{L}(\varphi))$ and subsequently checking $\mathcal{L}(\mathfrak{T})$ for emptiness. Since tree automata can be checked for emptiness in polynomial time [12, 7], this algorithm runs in exponential time in $|\varphi|$. As the problem of deciding VLDL satisfiability is EXPTIME-hard [23], the obtained algorithm is asymptotically optimal.

We split the proof of Theorem 2 into two parts: First, we transform a given VLDL formula into an equivalent so-called 1-AJA [5] of polynomial size. A 1-AJA is an alternating finite-state automaton on words that is able to “jump” from calls to their matching return, skipping the nested infix. We describe

this construction in the proof of Lemma 1. In a second step, we transform the obtained 1-AJA into a tree automaton of exponential size that recognizes the stack trees of words recognized by the 1-AJA. We define this construction in Lemma 3.

Let us first define the above mentioned 1-AJA [5]. First, let $Dirs = \{\rightarrow, \curvearrowright\}$. In general, we use \rightsquigarrow to indicate an arbitrary member of $Dirs$. Moreover, for a finite set Q and $\rightsquigarrow \in Dirs$, let $Comms_{\rightsquigarrow}(Q) = \{\rightsquigarrow\} \times Q \times Q$, let $Comms(Q) = Comms_{\rightarrow}(Q) \cup Comms_{\curvearrowright}(Q)$, and let $\mathcal{B}^+(Comms(Q))$ be the set of positive Boolean formulas over $Comms(Q)$. Note that $\mathcal{B}^+(Comms(Q))$ does not include the shorthands *true* nor *false*. A 1-AJA (with Büchi acceptance) $\mathfrak{A} = (Q, \tilde{\Sigma}, \delta, q_I, Q_F)$ consists of a finite set of states Q , a visibly pushdown alphabet $\tilde{\Sigma}$, a transition function $\delta: Q \times \Sigma \rightarrow \mathcal{B}^+(Comms(Q))$, an initial state $q_I \in Q$, and a set of accepting states $Q_F \subseteq Q$. We define $|\mathfrak{A}| = |Q|$.

Intuitively, when the automaton is in state q at position i of the word $\alpha = \alpha_0\alpha_1\alpha_2\cdots$, it guesses a set of commands $C \subseteq Comms(Q)$ such that $C \models \delta(q, \alpha_i)$. It then spawns one copy of itself for each command $(\rightsquigarrow, q_{\rightarrow}, q_{\curvearrowright}) \in C$ and executes the command with that copy. If $\rightsquigarrow = \curvearrowright$ and if α_i is a matched call, the copy jumps to the position of the matching return of α_i and transitions to state q_{\curvearrowright} . Otherwise, i.e., if $\rightsquigarrow = \rightarrow$, or if α_i is not a matched call, the automaton advances to position $i + 1$ and transitions to state q_{\rightarrow} . We say that \mathfrak{A} takes a jumping (direct) transition in the former (latter) case. All copies of \mathfrak{A} proceed in parallel. A single copy of \mathfrak{A} accepts if it visits accepting states infinitely often, while the 1-AJA accepts α if all of its copies accept.

Formally, a run of \mathfrak{A} on an infinite word $\alpha = \alpha_0\alpha_1\alpha_2\cdots$ is an infinite directed acyclic graph $R = (V, E)$ with $V \subseteq \mathbb{N} \times Q$, where $v_I = (0, q_I) \in V$ and all $v \in V$ are reachable from v_I . We call v_I the initial vertex of R and say that a vertex $(i, q) \in V$ is on level i of R . We require that for each $(i, q) \in V$, there exists some $C \subseteq Comms(Q)$ such that $C \models \delta(\alpha_i, q)$ and such that $((i, q), (i', q')) \in E$ if and only if $(i', q') = app(i, c)$ for some $c \in C$. To this end, the command-application function app is defined as $app(i, (\rightsquigarrow, q_{\rightarrow}, q_{\curvearrowright})) = (j, q_{\curvearrowright})$ if $\rightsquigarrow = \curvearrowright$ and α_i is a matched call with α_j as its matching return, and $app(i, (\rightsquigarrow, q_{\rightarrow}, q_{\curvearrowright})) = (i + 1, q_{\rightarrow})$ otherwise. We say that a vertex (i, q) is accepting if q is accepting. Furthermore, a run R is accepting if each vertex in R has at least one successor and if all infinite paths through R starting in v_I contain infinitely many accepting vertices.

Note that, in contrast to the classical definition of runs of alternating automata without jumping capability, an edge in the run of a 1-AJA does not characterize an advance by a single symbol. Instead, there exists “long” edges that characterize the automaton “jumping over” a nested infix. Thus, there may exist positions k such that a run of a 1-AJA on a word does not contain any vertices of the form (k, q) , since all copies of the automaton jump over position k . The cardinal positions of a word α , however, serve as synchronization points of a run on α , as no copy of the automaton is able to jump over the cardinal points.

Lemma 1. *For every VLDL formula φ there exists an effectively constructible 1-AJA \mathfrak{A} with $\mathcal{L}(\mathfrak{A}) = \mathcal{L}(\varphi)$ and with $|\mathfrak{A}| \in \mathcal{O}(p(|\varphi|))$ for some polynomial p .*

Proof. In earlier work, we constructed a 1-AJA with a more complicated condition from a given VLDL formula by induction over its structure [23]. This more complicated condition allowed for a complementation without a state-space-blowup in the construction of an automaton equivalent to $\varphi = \neg\varphi'$. As we are now aiming for a 1-AJA with a simpler acceptance condition, namely a Büchi-condition, we adapt this previous construction.

In order to prevent the costly complementation of 1-AJA, we require φ to be in negation normal form (NNF), i.e., we assume that negations only occur directly preceding atomic propositions. Should this not be the case, we can easily transform φ into NNF by “pushing down” negations along the syntax tree, using De Morgan’s law and the duality $\neg(\mathfrak{A})\varphi = [\mathfrak{A}]\neg\varphi$. Note that this latter duality does not require complementation of \mathfrak{A} , hence it is applicable in constant time. We then construct \mathfrak{A}_{φ} inductively over the structure of φ .

If $\varphi = p$, $\varphi = \neg p$, or $\varphi = \varphi_1 \circ \varphi_2$ for $\circ \in \{\vee, \wedge\}$, we trivially obtain \mathfrak{A}_{φ} , with $|\mathfrak{A}_p| = |\mathfrak{A}_{\neg p}| = 2$ and $|\mathfrak{A}_{\varphi_1 \circ \varphi_2}| \in \mathcal{O}(|\mathfrak{A}_{\varphi_1}| + |\mathfrak{A}_{\varphi_2}|)$ due to closure of 1-AJA under these operations [5]. If $\varphi = (\mathfrak{A})\varphi'$, we follow the same intuition as in the previous construction [23], i.e., we construct the 1-AJA \mathfrak{A}_{φ} such that a single copy jumps along the cardinal positions of the input-word and spawns copies at every matched call in order to verify that the jumps taken correctly summarize finite runs of \mathfrak{A} on the nested infix. Additionally, \mathfrak{A}_{φ} spawns copies verifying that the tests annotating the states along the simulated run hold true. Finally, \mathfrak{A}_{φ} nondeterministically decides to transition into $\mathfrak{A}_{\varphi'}$. The complete construction

for this case can be found in the full version of our previous work [23], which can be adapted to use a Büchi-condition by making none of the states simulating \mathfrak{A} accepting in \mathfrak{A}_φ , thus forcing the simulated run to eventually transition into $\mathfrak{A}_{\varphi'}$.

If $\varphi = [\mathfrak{A}]\varphi'$, we obtain an automaton equivalent to φ via a dual construction to the one described above for the case $\varphi = \langle \mathfrak{A} \rangle \varphi'$. Again, let $\mathfrak{A} = (Q^{\mathfrak{A}}, \tilde{\Sigma}, \Gamma^{\mathfrak{A}}, \Delta^{\mathfrak{A}}, q_I^{\mathfrak{A}}, Q_F^{\mathfrak{A}}, t^{\mathfrak{A}})$. By induction, we obtain the 1-AJAs $\mathfrak{A}' = (Q', \tilde{\Sigma}, \delta', q_I', Q_F')$ equivalent to φ' and, for each test φ_i occurring in \mathfrak{A} , let $\mathfrak{A}_i = (Q^i, \tilde{\Sigma}, \delta^i, q_I^i, Q_F^i)$ be the 1-AJA equivalent to $\neg\varphi_i$. Recall that we first transform $\neg\varphi_i$ into negation normal form by “pushing down” the negation to only occur in front of atomic propositions.

We now construct a 1-AJA \mathfrak{A}_φ equivalent to φ . This construction is dual to our previous one for the case $\varphi = \langle \mathfrak{A} \rangle \varphi'$ [23]. Intuitively, we simulate all runs of \mathfrak{A} on the input word by spawning copies of the main automaton that jump along the cardinal positions of the processed input word. Every time a call c is processed, we have to consider both cases of it being matched or unmatched. If c is matched, for each state $q \in Q^{\mathfrak{A}}$, we nondeterministically guess whether or not the automaton \mathfrak{A} can be in state q at the next step, or whether all run infixes starting in the current state lead to some state other than q . If c is unmatched, however, we treat c similarly to a local action and additionally denote that no unmatched return may be read anymore, since doing so would contradict c being unmatched. Note that, since 1-AJA process the matching return of a call after a jump instead of processing the symbol following it, we introduce waiting states that are used to delay execution for a single step.

Formally, we define the set of states $Q = Q_\varphi \cup \{\top\} \cup Q' \cup \bigcup_{\varphi_i \in \text{range}(t^{\mathfrak{A}})} Q^i$, where $Q_\varphi = \{q, q_{wait} \mid q \in (Q^{\mathfrak{A}} \times \{0, 1\}) \cup (Q^{\mathfrak{A}} \times Q^{\mathfrak{A}} \times \Gamma)\}$ and where the state \top is used as an accepting sink. The states from $Q^{\mathfrak{A}} \times \{0, 1\}$ are used to simulate the original automaton before ($Q^{\mathfrak{A}} \times \{0\}$) and after ($Q^{\mathfrak{A}} \times \{1\}$) processing at least one unmatched call.

For the sake of readability, we define the transition function for the different components of the automaton separately. Moreover, we write (\rightarrow, q) and (\rightarrow_a, q) as shorthands for (\rightarrow, q, \top) and (\rightarrow_a, \top, q) . As \top is used as a sink, we clearly have $\delta_{sink}(\top, a) = (\rightarrow, \top)$ for all $a \in \Sigma$. Furthermore, we define $\delta_{wait}(q_{wait}, a) = q$ for all $a \in \Sigma$.

When encountering a final state of \mathfrak{A} , we model acceptance of the prefix processed so far by spawning a copy that moves to the initial state of \mathfrak{A}' . To achieve a uniform presentation, we define the auxiliary formula $\chi^f(q, a) = \delta'(q_I', a)$ if $q \in Q_F^{\mathfrak{A}}$ and $\chi^f(q, a) = (\rightarrow, acc)$ otherwise.

Moreover, we need notation to denote transitions into the automata \mathfrak{A}_i implementing the negated tests of \mathfrak{A} . Note that, in order to handle the test labeling the initial state of \mathfrak{A} correctly, we only enter the automaton implementing the negation of a test upon leaving the respective state. Hence, we move to the successors of the initial state of \mathfrak{A}_i instead of moving to the initial state itself. To this end, we define the auxiliary formula $\theta_q^a = \delta^i(q_I^i, a)$, where $t(q) = \varphi_i$.

Upon reading a local action, we have to spawn a copy to continue in the automaton \mathfrak{A}' if we are currently in a final state, as well as copies to track all possible continuations of the subsequent run. If the test $t^{\mathfrak{A}}(q)$ is violated, however, we verify that this is indeed the case by moving to the automaton implementing $\neg t^{\mathfrak{A}}(q)$. Hence we have

$$\delta_{main}((q, b), l) = \left[\chi^f(q, l) \wedge \bigwedge_{(q, l, q') \in \Delta} (\rightarrow, (q', b)) \right] \vee \theta_q^l \text{ for } l \in \Sigma_l, b \in \{0, 1\}$$

Upon reading a call, the constructed 1-AJA has to consider both the case that the call is matched as well as that it is unmatched. In the former case, for all transitions $(q, c, q', A) \in \Delta$ and all states $q'' \in Q$, the automaton either spawns a copy that verifies that it is impossible to go from q' to q'' by popping A off the stack in the final transition, or it continues at the matching return in state q'' . In the latter case it ignores the effects on the stack and denotes that it may not read any returns from this point onwards by setting the binary flag in its state to 1. Similarly to the previous case, we can instead verify that $t(q)$ is violated by moving to the automaton implementing $\neg t(q)$.

$$\delta_{main}((q, b), c) = \left[\chi^f(q) \wedge \bigwedge_{(q, c, q', A) \in \Delta, q'' \in Q} \left[(\rightarrow, (q', q'', A)) \vee (\curvearrowright, (q'', b)_{wait}) \right] \wedge \bigwedge_{(q, c, q', A) \in \Delta} (\rightarrow, (q', 1)) \right] \vee \theta_q^c \text{ for } c \in \Sigma_c, b \in \{0, 1\}$$

The main automaton may only handle returns as long as it has not skipped any calls. If it encounters a return after having guessed that a call is unmatched, it moves to the accepting sink in order to be able

to continue the simulation of all remaining runs.

$$\begin{aligned}\delta_{main}((q, 0), r) &= \left[\chi^f(q, r) \wedge \bigwedge_{(q, r, \perp, q') \in \Delta} (\rightarrow, (q', 0)) \right] \vee \theta_q^r \quad \text{for } r \in \Sigma_r \\ \delta_{main}((q, 1), r) &= (\rightarrow, \top) \quad \text{for } r \in \Sigma_r\end{aligned}$$

The transition function δ_{main} determines the behavior of the main automaton. It remains to define the behavior of the copies of the automaton verifying the inability of the automaton to move to some particular state upon reading a matching return. These behave similarly to the main automaton on reading local actions and calls. The main difference in handling calls is that these automata do not need to guess whether or not a call is matched: Since they are only spawned on reading supposedly matched calls and terminate their run upon reading the matching return, all calls they encounter can be assumed to be matched as well. Additionally, they never transition to the automaton \mathfrak{A}' , but merely to the automaton implementing the negation of the test of the current state upon having verified their guess. If instead they moved, say, to the accepting sink \top , there would indeed be a possibility to move to the chosen state upon popping the given stack symbol, which would contradict the nondeterministic guess made upon reading the matching call.

$$\begin{aligned}\delta_{ver}((q, q', A), l) &= \left[\bigwedge_{(q, l, q'') \in \Delta} (\rightarrow, (q'', q', A)) \right] \vee \theta_q^l \quad \text{if } l \in \Sigma_l \\ \delta_{ver}((q, q', A), c) &= \left[\bigwedge_{(q, c, q'', A') \in \Delta, q''' \in Q} (\rightarrow, (q'', q''', A')) \vee (\neg, (q''', q', A)_{wait}) \right] \vee \theta_q^c \quad \text{if } c \in \Sigma_c \\ \delta_{ver}((q, q', A), r) &= \theta_q^r \quad \text{if } r \in \Sigma_r, (q, r, A, q') \in \Delta \\ \delta_{ver}((q, q', A), r) &= (\rightarrow, rej) \quad \text{if } r \in \Sigma_r, (q, r, A, q') \notin \Delta\end{aligned}$$

We then define the complete transition function δ of \mathfrak{A}_φ as the union of the previously defined partial transition functions. Since their domains are pairwise disjoint, this union is well-defined.

$$\delta = \delta_{sink} \cup \delta_{wait} \cup \delta' \cup \bigcup_{\varphi_i \in \text{range}(t)} \delta^i \cup \delta_{main} \cup \delta_{ver}$$

Finally, we make all states obtained by the translation of \mathfrak{A} accepting. Thus, the simulations of all runs of \mathfrak{A} are accepting, which lets the complete automaton \mathfrak{A}_φ track all prefixes of the processed word that are accepted by \mathfrak{A} . The 1-AJA

$$\mathfrak{A}_\varphi = (Q, \tilde{\Sigma}, \delta, (q_I^{\mathfrak{A}}, 0), Q_F \cup Q'_F \cup \bigcup_{\varphi_i \in \text{range}(t^{\mathfrak{A}})} Q_F^i)$$

then recognizes the language of $\varphi = \langle \mathfrak{A} \rangle \varphi'$, where $Q_F = (Q^{\mathfrak{A}} \times \{0, 1\}) \cup (Q^{\mathfrak{A}} \times Q^{\mathfrak{A}} \times \Gamma) \cup \{\top\}$. \square

Having given a translation of VLDL formulas into 1-AJAs, we now show how to transform a given 1-AJA \mathfrak{A} into a tree automaton recognizing the stack trees of words recognized by \mathfrak{A} . To this end, consider a run R of a 1-AJA \mathfrak{A} on some word $\alpha \in \Sigma^\omega$, as illustrated on the left-hand side of Figure 4. As argued above, the cardinal positions of the processed word serve as synchronization points in the run of \mathfrak{A} on α : If i is a cardinal position of α , then there exist no positions $j, j' \in \mathbb{N}$ with $j < i < j'$ such that R contains an edge from level j to level j' . In other words, each infinite path starting in the initial vertex v_I of R contains a vertex on level i for each cardinal position i of α . Hence, we are able to decide whether or not R is accepting by considering finite paths of R starting and ending in levels i and i' , respectively, where i and i' are cardinal positions of α .

More formally, we demonstrate that the breakpoint construction of Miyano and Hayashi [15] can be adapted to 1-AJAs. To this end, let $R = (V, E)$ and let V_i be the vertices occurring in R on level i , i.e., $V_i = \{(i, q) \in V \mid q \in Q\}$. A breakpoint sequence over R is an infinite sequence of cardinal positions $0 = i_0 < i_1 < i_2 \dots$ of α such that all finite paths in R starting on level i_j and ending on level i_{j+1} contain at least one accepting vertex. Each cardinal position i_j in a breakpoint sequence is called a breakpoint.

Lemma 2. *Let R be a run of a 1-AJA. The run R is accepting if and only if there exists a breakpoint sequence over R .*

Proof. First assume that there exists a breakpoint sequence $0 = i_0, i_1, i_2, \dots$ over R . Then R is clearly accepting, as each infinite path π starting in v_I is of the form $\pi = v_0\pi_0v_1\pi_1v_2\pi_2\cdots$, where each $v_j\pi_jv_{j+1}$ is a path from level i_j to level i_{j+1} , hence $v_j\pi_jv_{j+1}$ contains at least one accepting vertex. Thus, π is accepting.

For the other direction, assume that R is accepting. We show the existence of a breakpoint sequence inductively and begin by defining $i_0 = 0$. Now let i_0, \dots, i_j be a finite prefix of a breakpoint sequence and assume towards a contradiction that no cardinal position i_{j+1} exists such that i_0, \dots, i_j, i_{j+1} is a prefix of a breakpoint sequence. Then, for each cardinal position k , there exists a path from some vertex on level i_j to some vertex on level k that does not contain an accepting vertex. Hence, there also exists an infinite path starting on level i_j that does not contain an accepting vertex, which contradicts R being accepting. Thus, there exists a cardinal position i_{j+1} such that i_0, \dots, i_j, i_{j+1} is a prefix of some breakpoint sequence. Hence, there exists a breakpoint sequence over R . \square

Given some 1-AJA \mathfrak{A} , we now construct a tree automaton that verifies that the input tree is indeed a stack tree and, if this is the case, simulates a run of \mathfrak{A} on the word represented by the input tree by keeping track of the set of states at each level. Moreover, it verifies the existence of a breakpoint sequence, visiting an accepting state on the cardinal branch of the processed tree every time the corresponding symbol is at a cardinal position of the input word that can continue the prefix of the breakpoint sequence constructed so far. In order to do so, we adapt the breakpoint construction by Miyano and Hayashi [15]. The key insight of this construction is that, given some breakpoint i , the vertices of any level $j > i$ can be partitioned into two sets A_j and N_j . The set A_j contains those states such that each finite path from some vertex on level i to some vertex on level j visits at least one accepting state, while N_j contains the remaining vertices on level j . We illustrate this partitioning on the left-hand side of Figure 4. If the set N_j is empty, then j continues the breakpoint sequence constructed so far. We adapt this technique in order to translate 1-AJA into tree automata by keeping track of the sets A_i and N_i along the cardinal branch of the stack tree. Upon encountering a matched call at position i , the tree automaton guesses the sets A_j and N_j reached at the next cardinal position j and verifies this guess when processing the nested infix of position i .

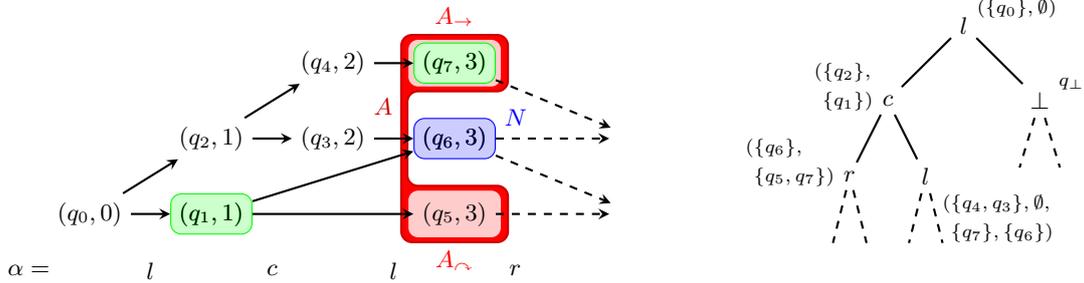


Fig. 4. Encoding of a run of a 1-AJA (left) into a run of a tree automaton (right). The states q_1 and q_7 are accepting. The positions 0, 1, and 3 are cardinal positions of α . Note $N = N_{\rightarrow} = N_{\sim}$.

Lemma 3. *For every 1-AJA \mathfrak{A} there exists an effectively constructible tree automaton \mathfrak{T} with $\mathcal{L}(\mathfrak{T}) = \text{st}(\mathcal{L}(\mathfrak{A}))$ and $|\mathfrak{T}| \in \mathcal{O}(2^{|\varphi|})$.*

Proof. We construct a tree automaton \mathfrak{T}' such that $\mathcal{L}(\mathfrak{T}') \cap \text{st}(\Sigma^\omega) = \text{st}(\mathcal{L}(\mathfrak{A}))$. Recall that, due to Theorem 1, we obtain a tree automaton \mathfrak{T}_Σ with $\mathcal{L}(\mathfrak{T}_\Sigma) = \text{st}(\Sigma^\omega)$. By intersecting \mathfrak{T}' with \mathfrak{T}_Σ we subsequently obtain \mathfrak{T} with the properties stated above.

We have explained the behavior of the automaton \mathfrak{T}' along the cardinal branch above. It remains to take into account the effect of nested infixes on the states reached by \mathfrak{A} at cardinal positions. To this end, we note that each state reached at a cardinal position is either reached by taking a jumping transition from the previous cardinal position, or by taking a direct transition from the directly preceding position, i.e., from the last position of the nested infix. Thus, when reading a matched call at position i , the automaton \mathfrak{T} guesses sets $A_{\rightarrow}, N_{\rightarrow} \subseteq Q$ that are reached eventually by copies of the automaton

that process the nested infix w of position i . It then assumes that these states are indeed reached by processing w and verifies that guess while processing $\text{st}(w)$, i.e., the right-hand subtree of the matched call. We show an example of this encoding of a run of \mathfrak{A} as a run of \mathfrak{T}' on the right-hand side of Figure 4.

We use two kinds of states in order to implement this idea. States of the form (A, N) , where A and N partition a nonempty subset of Q are used along the cardinal branch of the processed stack tree, implementing the breakpoint construction. Furthermore, we use states of the form $((A, N), (A_G, N_G))$, where (A, N) as well as (A_G, N_G) are partitions of nonempty subsets of Q , to verify guesses about the effects of processing nested infixes. Moreover, we use a sink-state q_\perp in order to process subtrees labelled exclusively with \perp .

Let 1-AJA $\mathfrak{A} = (Q, \tilde{\Sigma}, \delta, q_I, Q_F)$. In order to define \mathfrak{T}' concisely, we introduce some notation. First, let $\overline{Q_F} = Q \setminus Q_F$. Moreover, given some set $C \subseteq \text{Comms}(Q)$, we extract the direct- and jump-target-states $Q_{\rightsquigarrow}^C = \{q_{\rightsquigarrow} \mid (\rightsquigarrow', q_{\rightarrow}, q_{\curvearrowright}) \in C\}$ for $\rightsquigarrow \in \text{Dirs}$. Furthermore, given some nonempty set $B = \{\varphi_0, \dots, \varphi_n\}$ of Boolean formulas over $\text{Comms}(Q)$, we write $C \models B$ if C is a minimal model of B , i.e., if $C = \bigcup_{i=0}^n C_i$ with $C_i \models \varphi_i$.

Formally, we define the tree automaton $\mathfrak{T}' = (Q', \Sigma, \Delta, q'_I, Q'_F)$, where $Q' = (2^Q)^2 \cup ((2^Q)^2 \times (2^Q)^2) \cup \{q_\perp\}$, $q'_I = (\{q_I\}, \emptyset)$, if $q_I \in F$ and $q'_I = (\emptyset, \{q_I\})$ otherwise, $Q'_F = (2^Q \times \{\emptyset\}) \cup ((2^Q)^2 \times (2^Q)^2) \cup \{q_\perp\}$, and the transition relation Δ defined as the smallest relation that satisfies all of the following conditions:

Local Actions and Returns Let $x \in \Sigma_l \cup \Sigma_r$, let S be a nonempty subset of Q , and let (A, N) be a partition of S . Moreover, let $C_A, C_N \subseteq \text{Comms}(Q)$ such that $C_A \models \{\delta(q, x) \mid q \in A\}$ and such that $C_N \models \{\delta(q, x) \mid q \in N\}$. Since \mathfrak{A} cannot take jumping transitions upon processing local actions or returns, we define $N' = Q_{\rightarrow}^{C_N} \setminus F$ and $A' = (Q_{\rightarrow}^{C_A} \cap F) \cup (Q_{\rightarrow}^{C_A} \setminus N')$, thus updating the partitions A and N as described above. Note that the successor of a state in A may be in the subset N' if the same state is a successor of a state in N . Moreover, it is easy to verify that (A', N') indeed are a partition of some set $S' \subseteq Q$. Let (A_G, N_G) be a partition of some $S_G \subseteq Q$. We require $((A, N), (A_G, N_G)), x, ((A', N'), (A_G, N_G)), q_\perp \in \Delta$. Furthermore, if $N \neq \emptyset$, we require $((A, N), x, (A', N'), q_\perp) \in \Delta$.

Unmatched Calls Let $c \in \Sigma_c$, let S be a nonempty subset of Q , and let (A, N) be a partition of S . Moreover, let $C_A, C_N \subseteq \text{Comms}(Q)$ such that $C_A \models \{\delta(q, x) \mid q \in A\}$ and such that $C_N \models \{\delta(q, x) \mid q \in N\}$. In this case, we guess that the currently processed call is unmatched. Thus, similarly to the previous case, \mathfrak{A} can only take direct transitions. Define $N' = Q_{\rightarrow}^{C_N} \setminus F$ and $A' = (Q_{\rightarrow}^{C_A} \cap F) \cup (Q_{\rightarrow}^{C_A} \setminus N')$. We require $((A, N), c, q_\perp, (A', N')) \in \Delta$. We do, however, not require a transition processing unmatched calls when verifying some guess along a nested infix, as unmatched calls cannot occur in nested infixes.

Matched Calls Let $c \in \Sigma_c$ let $S \subseteq Q$ be nonempty, and let (A, N) be a partition of S . Moreover, let $C_A^{\rightsquigarrow}, C_N^{\rightsquigarrow} \subseteq \text{Comms}_{\rightsquigarrow}(Q)$ for $\rightsquigarrow \in \text{Dirs}$ such that $(C_A^{\rightarrow} \cup C_A^{\curvearrowright}) \models \{\delta(q, x) \mid q \in A\}$ and such that $(C_N^{\rightarrow} \cup C_N^{\curvearrowright}) \models \{\delta(q, x) \mid q \in N\}$. We follow the same idea as in the previous two cases and first define the sets of states reached by \mathfrak{A} at the first position of the nested infix. To this end, let $N'_{\rightarrow} = Q_{\rightarrow}^{C_N^{\rightsquigarrow}} \setminus F$ and let $A'_{\rightarrow} = (Q_{\rightarrow}^{C_A^{\rightsquigarrow}} \setminus N'_{\rightarrow}) \cup (Q_{\rightarrow}^{C_N^{\rightsquigarrow}} \cap F)$. Moreover, we define the partition of states reached by \mathfrak{A} at the matching return of the current letter by taking a jumping transition as $N'_{\curvearrowright} = (Q_{\curvearrowright}^{C_N^{\rightsquigarrow}} \setminus F)$ and $A'_{\curvearrowright} = (Q_{\curvearrowright}^{C_A^{\rightsquigarrow}} \setminus N'_{\curvearrowright}) \cup (Q_{\curvearrowright}^{C_N^{\rightsquigarrow}} \cap F)$. Finally, we guess that \mathfrak{A} finishes processing the nested infix with the partition (A_G, N_G) of some nonempty $S_G \subseteq Q$. Thus, we require $((A, N, A'_G, N'_G), c, (A'_{\curvearrowright} \cup A_G, N'_{\curvearrowright} \cup N_G), A'_G, N'_G), (A'_{\rightarrow}, N'_{\rightarrow}, A_G, N_G) \in \Delta$ for arbitrary nonempty $S'_G \subseteq Q$ where (A'_G, N'_G) is a partition of S'_G . Moreover, if $N \neq \emptyset$, we require $((A, N), c, (A'_{\curvearrowright}, N'_{\curvearrowright}), (A'_{\rightarrow}, N'_{\rightarrow}, A_G, N_G)) \in \Delta$.

Breakpoint Let $x \in \Sigma$, let $S \subseteq Q$ be nonempty, and let $A = S \cap F$ and $N = S \setminus F$. If $((A, N), x, q'_0, q'_1) \in \Delta$, we require $((S, \emptyset), x, q'_0, q'_1) \in \Delta$.

Verified Guess Let $S \subseteq Q$ be nonempty and let (A, N) be a partition of S . We require $((A, N, A, N), \perp, q_\perp, q_\perp) \in \Delta$.

Sink We require $(q_\perp, \perp, q_\perp, q_\perp) \in \Delta$.

Let α be some word accepted by \mathfrak{A} and let $t = \text{st}(\alpha)$. Then one can easily construct a run of \mathfrak{T}' on t from a run of \mathfrak{A} on α as indicated in Figure 4. In fact, if there exists an accepting run of \mathfrak{A} on α , then there also exists an accepting run of \mathfrak{T}' on t , due to the implementation of the breakpoint construction and due to Lemma 2. Conversely, if \mathfrak{T}' accepts a stack tree $t = \text{st}(\alpha)$ for some word $\alpha \in \Sigma^\omega$, say with the accepting run R' , then it is possible to reconstruct an accepting run R of \mathfrak{A} on α from R' via a preorder-traversal of R' that traverses the right-hand children of vertices first. Again, due to Lemma 2,

the run R is accepting if and only if R' is accepting. Hence, by intersecting \mathfrak{T}' with the automaton \mathfrak{T}_Σ recognizing $\text{st}(\Sigma^\omega)$ we obtain the automaton \mathfrak{T} recognizing $\text{st}(\mathcal{L}(\mathfrak{A}))$. As $|\mathfrak{T}'| \in \mathcal{O}(2^{|\mathfrak{A}|})$, and since \mathfrak{T}_Σ is of fixed size, we obtain $|\mathfrak{T}| \in \mathcal{O}(2^{|\mathfrak{A}|})$. \square

The proof of Theorem 2 follows from Lemma 1 and Lemma 3: Given a VLDL formula φ , we first construct the 1-AJA \mathfrak{A} with $\mathcal{L}(\mathfrak{A}) = \mathcal{L}(\varphi)$ as demonstrated in the proof of Lemma 1. The automaton \mathfrak{A} is of size polynomial in $|\varphi|$. We then construct the tree automaton \mathfrak{T} with $\mathcal{L}(\mathfrak{T}) = \text{st}(\mathcal{L}(\mathfrak{A}))$ as shown in the proof of Lemma 3. The automaton \mathfrak{T} recognizes $\text{st}(\mathcal{L}(\mathfrak{A})) = \text{st}(\mathcal{L}(\varphi))$ and is of size exponential in $|\mathfrak{A}|$, i.e., of size exponential in $|\varphi|$.

5 Reducing VLDL Model Checking to Tree Automata Emptiness

In the previous section we have reduced the problem of VLDL satisfiability checking to the emptiness problem for tree automata. We now consider the problem of VLDL model checking, which is formulated as follows: “Given a VPS \mathcal{S} and a VLDL formula φ , does $\text{traces}(\mathcal{S}) \subseteq \mathcal{L}(\varphi)$ hold true?” We now show that this problem can be reduced to the emptiness problem for tree automata similarly to the reduction of the satisfiability problem for VLDL to the same problem.

Theorem 3. *Let \mathcal{S} be a VPS and let φ be a VLDL formula. There exists an effectively constructible tree automaton \mathfrak{T} such that $\mathcal{L}(\mathfrak{T}) = \emptyset$ if and only if $\text{traces}(\mathcal{S}) \subseteq \mathcal{L}(\varphi)$ with $|\mathfrak{T}| \in \mathcal{O}(2^{|\varphi|}p(|\mathcal{S}|))$ for some polynomial p .*

Proof. Recall that $\text{traces}(\mathcal{S}) \subseteq \mathcal{L}(\varphi)$ if and only if $\text{traces}(\mathcal{S}) \cap \mathcal{L}(\neg\varphi) = \emptyset$. Moreover, recall that we can effectively construct a tree automaton $\mathfrak{T}_{\neg\varphi}$ such that $\mathcal{L}(\mathfrak{T}_{\neg\varphi}) = \text{st}(\mathcal{L}(\neg\varphi))$ due to Theorem 2. We now construct a tree automaton $\mathfrak{T}_\mathcal{S}$ recognizing $\text{st}(\text{traces}(\mathcal{S}))$. By intersecting $\mathfrak{T}_\mathcal{S}$ and $\mathfrak{T}_{\neg\varphi}$ we subsequently obtain the tree automaton \mathfrak{T} recognizing $\text{traces}(\mathcal{S}) \cap \mathcal{L}(\neg\varphi)$. Hence, $\mathcal{L}(\mathfrak{T}) = \emptyset$ if and only if $\text{traces}(\mathcal{S}) \subseteq \mathcal{L}(\varphi)$.

It remains to construct $\mathfrak{T}_\mathcal{S}$. Similarly to the proof of Lemma 3, we first construct $\mathfrak{T}'_\mathcal{S}$ such that $\mathcal{L}(\mathfrak{T}'_\mathcal{S}) \cap \text{st}(\Sigma^\omega) = \text{st}(\text{traces}(\mathcal{S}))$. By intersecting $\mathfrak{T}'_\mathcal{S}$ with \mathfrak{T}_Σ we then obtain the required $\mathfrak{T}_\mathcal{S}$. The idea behind the construction of $\mathfrak{T}'_\mathcal{S}$ is to simulate a run of \mathcal{S} along the cardinal branch of the tree. This is straightforward in the case of local actions and unmatched calls or returns. Upon encountering a matched call, $\mathfrak{T}'_\mathcal{S}$ guesses the state reached by \mathcal{S} upon encountering the matched return and verifies that guess on the stack tree of the nested infix.

Let $\mathcal{S} = (Q, \tilde{\Sigma}, \Gamma, \Delta, q_I)$. We define $\mathfrak{T}'_\mathcal{S} = (Q', q_I, \Delta', Q'_F)$ with $Q' = Q \cup (Q \times Q) \cup (Q \times \Gamma) \cup (Q \times \Gamma \times Q) \cup \{q_\perp\}$, $Q'_F = Q'$, and

$$\Delta' = \Delta_l \cup \Delta_{uc} \cup \Delta_{ur} \cup \Delta_{mc} \cup \Delta_{mr} \cup \Delta_s \ .$$

The individual components of Δ' are defined as follows: We process local actions using transitions of the form

$$\Delta_l = \{(q, l, q', q_\perp), ((q, q_G), l, (q', q_G), q_\perp) \mid (q, l, q') \in \Delta, q_G \in Q\} \ .$$

Similarly, upon encountering unmatched calls or returns, we use transitions of the form

$$\Delta_{uc} = \{(q, c, q_\perp, q') \mid (q, c, q', A) \in \Delta\}$$

and

$$\Delta_{ur} = \{(q, r, q', q_\perp) \mid (q, r, \perp, q') \in \Delta\} \ ,$$

respectively. When encountering a matched call, we guess a state q_G reached by the automaton upon processing the matching return and verify that guess using transitions from

$$\Delta_{mc} = \{(q, c, (q_G, A), (q', q_G)) \mid (q, c, q', A) \in \Delta, q_G \in Q\} \cup \\ \{(q, c, q', A) \in \Delta, q_G, q'_G \in Q\} \ .$$

Upon encountering a matched return, we are in some state from $(Q \times \Gamma) \cup (Q \times \Gamma \times Q)$, since a matched return only occurs directly following a matched call. Hence, we use a transition from

$$\Delta_{mr} = \{((q, A), r, q', q_{\perp}) \mid (q, r, A, q') \in \Delta\} \cup \{((q, A, q_G), r, (q', q_G), q_{\perp}) \mid (q, r, A, q') \in \Delta\}$$

in order to process that matched return. Finally, we define $\Delta_s = \{(q_{\perp}, \perp, q_{\perp}, q_{\perp})\}$ to continue the run of $\mathfrak{T}'_{\mathcal{S}}$ upon encountering the sink state q_{\perp} .

Using the intuition given above, it can easily be verified that $\mathfrak{T}'_{\mathcal{S}} \cap \text{st}(\Sigma^{\omega}) = \text{st}(\text{traces}(\mathcal{S}))$ indeed holds true. Thus, as previously argued, we obtain the automaton $\mathfrak{T}_{\mathcal{S}, -\varphi}$ with the properties given in the statement of this lemma. \square

Due to Theorem 3, we obtain a novel asymptotically optimal algorithm for VLDL model checking: Given a VPS \mathcal{S} and a VLDL formula φ , we construct \mathfrak{T} such that $\mathcal{L}(\mathfrak{T}) = \emptyset$ if and only if $\text{traces}(\mathcal{S}) \subseteq \mathcal{L}(\varphi)$. The automaton \mathfrak{T} can be constructed in exponential time and is of exponential size in $|\varphi|$ and of polynomial size in $|\mathcal{S}|$. Hence, we can check \mathfrak{T} for emptiness in exponential time in $|\varphi|$ and in polynomial time in $|\mathcal{S}|$. Since the problem of VLDL model checking is EXPTIME-complete [23], this algorithm is asymptotically optimal.

6 Conclusion

In this work we have presented a correspondence between infinite words over a pushdown alphabet and infinite binary trees. Moreover, we demonstrated a construction translating VLDL formulas into tree automata that are language-equivalent with respect to the above correspondence. This construction yields novel algorithms for satisfiability- and model checking of VLDL formulas that reduce the problem to the emptiness problem for tree automata. Thus, this construction leverages the strong connection between visibly pushdown languages and regular tree languages that was already exhibited by Alur and Madhusudan in their seminal work on the former family of languages [2]. Moreover, the construction demonstrates that the well-known breakpoint construction by Miyano and Hayashi [15], which is routinely used to remove alternation from stack-free automata, can easily be adapted to transform alternating automata over visibly pushdown words into corresponding alternation-free automata over trees representing such words.

In future work, we plan to empirically evaluate both the algorithms presented in this work as well as those presented in earlier work [23], which reduce the satisfiability- and model checking problems for VLDL to the emptiness problem for visibly pushdown automata. Recall that our novel algorithm reduces both problems to the emptiness problem for tree automata, which in turn reduces to the well-studied problem of solving a two-player Büchi game. The latter problem is well-studied due to its important applications, e.g., in program verification [1, 22] and program synthesis [13]. Hence, there exist efficient algorithms [7] for solving them as well as mature solvers [10, 11]. Thus, we expect our novel algorithm to outperform the previous approach [23] to the above problems.

Moreover, in previous work we investigated the problem of solving two-player games on a visibly pushdown arena in which the winning condition is given by a VLDL formula and determined this problem to be 3EXPTIME-complete [23]. We showed membership of this problem in 3EXPTIME by reducing it to the problem of solving visibly pushdown games against a winning condition given by visibly pushdown automata. Currently, we are investigating whether there exists a reduction of the former problem to that of solving games in which the winning condition is given via tree automata that yields an asymptotically optimal algorithm.

Acknowledgements The author would like to thank Martin Zimmermann for multiple fruitful discussions.

References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *J. ACM* 49(5), 672–713 (2002)
2. Alur, R., Madhusudan, P.: Visibly pushdown languages. In: Babai, L. (ed.) *STOC 2004*. pp. 202–211. ACM (2004)
3. Ball, T., Rajamani, S.K.: *Bebop: A Symbolic Model Checker for Boolean Programs*. In: Havelund, K., Penix, J., Visser, W. (eds.) *SPIN 2000*. LNCS, vol. 1885, pp. 113–130. Springer (2000)

4. Ball, T., Rajamani, S.K.: Bebop: a path-sensitive interprocedural dataflow engine. In: Field, J., Snelting, G. (eds.) PASTE 2001. pp. 97–103. ACM (2001)
5. Bozzelli, L.: Alternating Automata and a Temporal Fixpoint Calculus for Visibly Pushdown Languages. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703 (2007)
6. Bozzelli, L., Sánchez, C.: Visibly Linear Temporal Logic. *J. Aut. Reas.* In Press
7. Chatterjee, K., Henzinger, M.: An $\mathcal{O}(n^2)$ time algorithm for alternating büchi games. In: Rabani, Y. (ed.) SODA 2012. pp. 1386–1399. SIAM (2012)
8. Chatterjee, K., Henzinger, T.A., Piterman, N.: Algorithms for Büchi Games. In: GDV (2006), available at <http://arxiv.org/abs/0805.2620>
9. Fijalkow, N., Pinchinat, S., Serre, O.: Emptiness Of Alternating Tree Automata Using Games With Imperfect Information. In: Seth, A., Vishnoi, N.K. (eds.) FSTTCS 2013. LIPIcs, vol. 24, pp. 299–311. Schloss Dagstuhl - LZI (2013)
10. Friedmann, O., Lange, M.: The PGSolver collection of parity game solvers. University of Munich (2009), available at github.com/tcsprojects/pgsolver/blob/master/doc/pgsolver.pdf.
11. Keiren, J.: An experimental study of algorithms and optimisations for parity games, with an application to Boolean Equation Systems. Master’s thesis, Eindhoven University of Technology (2009)
12. Kupferman, O., Vardi, M.Y.: Weak alternating automata and tree automata emptiness. In: Vitter, J.S. (ed.) STOC 1998. pp. 224–233. ACM (1998)
13. Kupferman, O., Vardi, M.Y.: From linear time to branching time. *ACM Trans. Comput. Log.* 6(2), 273–294 (2005)
14. Löding, C., Madhusudan, P., Serre, O.: Visibly pushdown games. In: Lodaya, L., Mahajan, M. (eds.) FSTTCS 2004. LNCS, vol. 3328, pp. 408–420. Springer (2005)
15. Miyano, S., Hayashi, T.: Alternating finite automata on ω -words. *TCS* 32(3), 321–330 (1984)
16. Nießner, F.: Nondeterministic Tree Automata. In: Grädel, E., Thomas, W., Wilke, T. (eds.) Automata, Logics, and Infinite Games: A Guide to Current Research. LNCS, vol. 2500, pp. 135–152. Springer (2001)
17. Pnueli, A.: The temporal logic of programs. In: FOCS 1977. pp. 46–57. IEEE (1977)
18. Schwoon, S.: Model checking pushdown systems. Ph.D. thesis, TU Munich (2002), Available at <http://tumb1.biblio.tu-muenchen.de/publ/diss/in/2002/schwoon.html>
19. Suwimonterabuth, D., Schwoon, S., Esparza, J.: jMoped: A Java Bytecode Checker Based on Moped. In: Halbwachs, N., Zuck, L.D. (eds.) TACAS 2005. LNCS, vol. 3440, pp. 541–545. Springer (2005)
20. Thomas, W.: Automata on infinite objects. *Handbook of theoretical computer science, Volume B* pp. 133–191 (1990)
21. Thomas, W.: Languages, automata, and logic. In: *Handbook of formal languages*, pp. 389–455. Springer (1997)
22. Vardi, M.Y.: Automata-theoretic model checking revisited. In: Chockler, H., Hu, A.J. (eds.) HVC 2008. LNCS, vol. 5394. Springer (2008)
23. Weinert, A., Zimmermann, M.: Visibly linear dynamic logic. In: Lal, A., Akshay, S., Saurabh, S., Sen, S. (eds.) FSTTCS 2016. LIPIcs, vol. 65, pp. 28:1–28:14. Schloss Dagstuhl - LZI (2016), full version available at <http://arxiv.org/abs/1512.05177>
24. Zielonka, W.: Infinite games on finitely coloured graphs with applications to automata on infinite trees. *TCS* 200(1-2), 135–183 (1998)