# Nested Antichains for WS1S

Tomáš Fiedor, Lukáš Holík, Ondřej Lengál, and Tomáš Vojnar

FIT, Brno University of Technology, IT4Innovations Centre of Excellence, Czech Republic

**Abstract.** We propose a novel approach for coping with alternating quantification as the main source of nonelementary complexity of deciding WS1S formulae. Our approach is applicable within the state-of-the-art automata-based WS1S decision procedure implemented, e.g., in MONA. The way in which the standard decision procedure processes quantifiers involves determinization, with its worst case exponential complexity, for every quantifier alternation in the prefix of a formula. Our algorithm avoids building the deterministic automata—instead, it constructs only those of their states needed for (dis)proving validity of the formula. It uses a symbolic representation of the states, which have a deeply nested structure stemming from the repeated implicit subset construction, and prunes the search space by a nested subsumption relation, a generalization of the one used by the so-called antichain algorithms for handling nondeterministic automata. We have obtained encouraging experimental results, in some cases outperforming MONA by several orders of magnitude.

## 1 Introduction

Weak monadic second-order logic of one successor (WS1S) is a powerful, concise, and decidable logic for describing regular properties of finite words. Despite its nonelementary worst case complexity [1], it has been shown useful in numerous applications. Most of the successful applications were due to the tool MONA [2], which implements a finite automata-based decision procedure for WS1S and WS2S (a generalization of WS1S to finite binary trees). The authors of MONA list a multitude of its diverse applications [3], ranging from software and hardware verification through controller synthesis to computational linguistics, and further on. Among more recent applications, verification of pointer programs and deciding related logics [4,5,6,7,8] can be mentioned, as well as synthesis from regular specifications [9]. MONA is still the standard tool and the most common choice when it comes to deciding WS1S/WS2S. There are other related automata-based tools that are more recent, such as jMosel [10] for a logic M2L(Str), and other than automata-based approaches, such as [11]. They implement optimizations that allow to outperform MONA on some benchmarks, however, none provides an evidence of being consistently more efficient. Despite many optimizations implemented in MONA and the other tools, the worst case complexity of the problem sometimes strikes back. Authors of methods using the translation of their problem to WS1S/WS2S are then forced to either find workarounds to circumvent the complexity blowup, such as in [5], or, often restricting the input of their approach, give up translating to WS1S/WS2S altogether [12].

The decision procedure of MONA works with deterministic automata; it uses determinization extensively and relies on minimization of deterministic automata to suppress the complexity blow-up. However, the worst case exponential complexity of determinization often significantly harms the performance of the tool. Recent works on

efficient methods for handling nondeterministic automata suggest a way of alleviating this problem, in particular works on efficient testing of language inclusion and universality of finite automata [13,14,15] and size reduction [16,22] based on a simulation relation. Handling nondeterministic automata using these methods, while avoiding determinization, has been shown to provide great efficiency improvements in [24] (abstract regular model checking) and also [23] (shape analysis). In this paper, we make a major step towards building the entire decision procedure of WS1S on nondeterministic automata using similar techniques. We propose a generalization of the antichain algorithms of [13] to that addresses the main bottleneck of the automata-based decision procedure for WS1S, which is also the source of its nonelementary complexity: elimination of alternating quantifiers on the automata level.

More concretely, the automata-based decision procedure translates the input WS1S formula into a finite word automaton such that its language represents exactly all models of the formula. The automaton is built in a bottom-up manner according to the structure of the formula, starting with predefined atomic automata for literals and applying a corresponding automata operation for every logical connective and quantifier ($\wedge, \vee, \neg, \exists$). The cause of the nonelementary complexity of the procedure can be explained on an example formula of the form $\varphi' = \exists X_m \forall X_{m-1} \ldots \forall X_2 \exists X_1 : \varphi_0$. The universal quantifiers are first replaced by negation and existential quantification, which results in $\varphi = \exists X_m \neg \exists X_{m-1} \ldots \neg \exists X_2 \neg \exists X_1 : \varphi_0$. The algorithm then builds a sequence of automata for the sub-formulae $\varphi_0, \varphi_0^\sharp, \ldots, \varphi_{m-1}, \varphi_{m-1}^\sharp$ of $\varphi$ where for $0 \leq i < m$, $\varphi_i^\sharp = \exists X_{i+1} : \varphi_i$, and $\varphi_{i+1} = \neg \varphi_i^\sharp$. Every automaton in the sequence is created form the previous one by applying the automata operations corresponding to negation or elimination of the existential quantifier, the latter of which may introduce nondeterminism. Negation applied on a nondeterministic automaton may then yield an exponential blowup: given an automaton for $\psi$, the automaton for $\neg \psi$ is constructed by the classical automata-theoretic construction consisting of determinization by the subset construction followed by swapping of the sets of final and non-final states. The subset construction is exponential in the worst case. The worst case complexity of the procedure run on $\varphi$ is then a tower of exponentials with one level for every quantifier alternation in $\varphi$; note that we cannot do much better—this non-elementary complexity is an inherent property of the problem.

Our new algorithm for processing alternating quantifiers in the prefix of a formula avoids the explicit determinization of automata in the classical procedure and significantly reduces the state space explosion associated with it. It is based on a generalization of the antichain principle used for deciding universality and language inclusion of finite automata [14,15]. It generalizes the antichain algorithms so that instead of being used to process only one level of the chain of automata, it processes the whole chain of quantifications with $i$ alternations on-the-fly. This leads to working with automata states that are sets of sets of sets ... of states of the automaton representing $\varphi_0$ of the nesting depth $i$ (this corresponds to $i$ levels of subset construction being done on-the-fly). The algorithm uses nested symbolic terms to represent sets of such automata states and a generalized version of antichain subsumption pruning which descends recursively down the structure of the terms while pruning on all its levels.

Our nested antichain algorithm can be in its current form used only to process a quantifier prefix of a formula, after which we return the answer to the validity query,

2

but not an automaton representing all models of the input formula. That is, we cannot use the optimized algorithm for processing inner negations and alternating quantifiers which are not a part of the quantifier prefix. However, despite this and the fact that our implementation is far less mature than that of MONA, our experimental results still show significant improvements over its performance, especially in terms of generated state space. We consider this a strong indication that using techniques for nondeterministic automata to decide WS1S (and WS$k$S) is highly promising. There are many more opportunities of improving the decision procedure based on nondeterministic automata, by using techniques such as simulation relations or bisimulation up-to congruence [17], and applying them to process not only the quantifier prefix, but all logical connectives of a formula. We consider this paper to be the first step towards a decision procedure for WS1S/WS$k$S with an entirely different scalability than the current state-of-the-art.

*Plan of the paper.* We define the logic WS1S in Section 2. In Sections 3 and 4, we introduce finite word automata and describe the classical decision procedure for WS1S based on finite word automata. In Section 5, we introduce our method for dealing with alternating quantifiers. Finally, we give an experimental evaluation and conclude the paper in Sections 6 and 7.

## 2 WS1S

In this section we introduce the *weak monadic second-order logic of one successor* (WS1S). We introduce only its minimal syntax here, for the full standard syntax and a more thorough introduction, see Section 3.3 in [18].

WS1S is a monadic second-order logic over the universe of discourse $\mathbb{N}_0$. This means that the logic allows second-order *variables*, usually denoted using upper-case letters $X, Y, \ldots$, that range over finite subsets of $\mathbb{N}_0$, e.g. $X = \{0, 3, 42\}$. Atomic formulae are of the form (i) $X \subseteq Y$, (ii) $\mathrm{Sing}(X)$, (iii) $X = \{0\}$, and (iv) $X = Y + 1$, where $X$ and $Y$ are variables. The atomic formulae are interpreted in turn as (i) standard set inclusion, (ii) the singleton predicate, (iii) $X$ is a singleton containing 0, and (iv) $X = \{x\}$ and $Y = \{y\}$ are singletons and $x$ is the successor of $y$, i.e. $x = y + 1$. Formulae are built from the atomic formulae using the logical connectives $\wedge, \vee, \neg$, and the quantifier $\exists X$ (for a second-order variable $X$).

Given a WS1S formula $\varphi(X_1, \ldots, X_n)$ with free variables $X_1, \ldots, X_n$, the assignment $\rho = \{X_1 \mapsto S_1, \ldots, X_n \mapsto S_n\}$, where $S_1, \ldots, S_n$ are finite subsets of $\mathbb{N}_0$, *satisfies* $\varphi$, written as $\rho \models \varphi$, if the formula holds when every variable $X_i$ is replaced with its corresponding value $S_i = \rho(X_i)$. We say that $\varphi$ is *valid*, denoted as $\models \varphi$, if it is satisfied by all assignments of its free variables to finite subsets of $\mathbb{N}_0$. Observe the limitation to *finite* subsets of $\mathbb{N}_0$ (related to the adjective *weak* in the name of the logic); a WS1S formula can indeed only have finite models (although there may be infinitely many of them).

## 3 Preliminaries and Finite Automata

For a set $D$ and a set $\mathbb{S} \subseteq 2^D$ we use $\downarrow\mathbb{S}$ to denote the *downward closure* of $\mathbb{S}$, i.e. $\downarrow\mathbb{S} = \{R \subseteq D \mid \exists S \in \mathbb{S} : R \subseteq S\}$, and $\uparrow\mathbb{S}$ to denote the *upward closure* of $\mathbb{S}$, i.e. $\uparrow\mathbb{S} = \{R \subseteq D \mid \exists S \in \mathbb{S} : R \supseteq S\}$. The set $\mathbb{S}$ is in both cases called the set of *generators* of $\uparrow\mathbb{S}$ or $\downarrow\mathbb{S}$, respectively. A set $\mathbb{S}$ is *downward closed* if it equals its downward closure, $\mathbb{S} = \downarrow\mathbb{S}$,

and *upward closed* if it equals to its upward closure, $\mathbb{S} = \uparrow\mathbb{S}$. The *choice* operator $\coprod$ (sometimes called the unordered Cartesian product) is an operator that, given a set of sets $\mathbb{D} = \{D_1, \ldots, D_n\}$, returns the set of all sets $\{d_1, \ldots, d_n\}$ obtained by taking one element $d_i$ from every set $D_i$. Formally,

$$\coprod\mathbb{D} = \big\{\{d_1, \ldots, d_n\} \mid (d_1, \ldots, d_n) \in \prod_{i=1}^{n} D_i\big\} \tag{1}$$

where $\prod$ denotes the Cartesian product. Note that for a set $D$, $\coprod\{D\}$ is the set of all singleton subsets of $D$, i.e. $\coprod\{D\} = \{\{d\} \mid d \in D\}$. Further note that if any $D_i$ is the empty set $\emptyset$, the result is $\coprod\mathbb{D} = \emptyset$.

Let $\mathbb{X}$ be a set of variables. A *symbol* $\tau$ over $\mathbb{X}$ is a mapping of all variables in $\mathbb{X}$ to either 0 or 1, e.g. $\tau = \{X_1 \mapsto 0, X_2 \mapsto 1\}$ for $\mathbb{X} = \{X_1, X_2\}$. An *alphabet* over $\mathbb{X}$ is the set of all symbols over $\mathbb{X}$, denoted as $\Sigma_{\mathbb{X}}$. For any $\mathbb{X}$ (even empty), we use $\overline{0}$ to denote the symbol which maps all variables from $\mathbb{X}$ to 0, $\overline{0} \in \Sigma_{\mathbb{X}}$.

A (nondeterministic) *finite* (word) *automaton* (abbreviated as FA) over a set of variables $\mathbb{X}$ is a quadruple $\mathcal{A} = (Q, \Delta, I, F)$ where $Q$ is a finite set of states, $I \subseteq Q$ is a set of *initial* states, $F \subseteq Q$ is a set of *final* states, and $\Delta$ is a set of transitions of the form $(p, \tau, q)$ where $p, q \in Q$ and $\tau \in \Sigma_{\mathbb{X}}$. We use $p \xrightarrow{\tau} q \in \Delta$ to denote that $(p, \tau, q) \in \Delta$. Note that for an FA $\mathcal{A}$ over $\mathbb{X} = \emptyset$, $\mathcal{A}$ is a unary FA with the alphabet $\Sigma_{\mathbb{X}} = \{\overline{0}\}$.

A *run* $r$ of $\mathcal{A}$ over a word $w = \tau_1\tau_2\ldots\tau_n \in \Sigma_{\mathbb{X}}^*$ from the state $p \in Q$ to the state $s \in Q$ is a sequence of states $r = q_0 q_1 \ldots q_n \in Q^+$ such that $q_0 = p$, $q_n = s$ and for all $1 \leq i \leq n$ there is a transition $q_{i-1} \xrightarrow{\tau_i} q_i$ in $\Delta$. If $s \in F$, we say that $r$ is an *accepting run*. We write $p \xRightarrow{w} s$ to denote that there exists a run from the state $p$ to the state $s$ over the word $w$. The *language* accepted by a state $q$ is defined by $\mathcal{L}_{\mathcal{A}}(q) = \{w \mid q \xRightarrow{w} q_f, q_f \in F\}$, while the language of a set of states $S \subseteq Q$ is defined as $\mathcal{L}_{\mathcal{A}}(S) = \bigcup_{q \in S} \mathcal{L}_{\mathcal{A}}(q)$. When it is clear which TA $\mathcal{A}$ we refer to, we only write $\mathcal{L}(q)$ or $\mathcal{L}(S)$. The language of $\mathcal{A}$ is defined as $\mathcal{L}(\mathcal{A}) = \mathcal{L}_{\mathcal{A}}(I)$. We say that the state $q$ accepts $w$ and that the automaton $\mathcal{A}$ accepts $w$ to express that $w \in \mathcal{L}_{\mathcal{A}}(q)$ and $w \in \mathcal{L}(\mathcal{A})$ respectively. We call a language $L \subseteq \Sigma_{\mathbb{X}}^*$ *universal* iff $L = \Sigma_{\mathbb{X}}^*$.

For a set of states $S \subseteq Q$, we define $post_{[\Delta,\tau]}(S) = \bigcup_{s \in S}\{t \mid s \xrightarrow{\tau} t \in \Delta\}$, $pre_{[\Delta,\tau]}(S) = \bigcup_{s \in S}\{t \mid t \xrightarrow{\tau} s \in \Delta\}$, and $cpre_{[\Delta,\tau]}(S) = \{t \mid post_{[\Delta,\tau]}(\{t\}) \subseteq S\}$.

The *complement* of $\mathcal{A}$ is the automaton $\mathcal{A}_{\mathcal{C}} = (2^Q, \Delta_{\mathcal{C}}, \{I\}, \downarrow\{Q \setminus F\})$ where $\Delta_{\mathcal{C}} = \big\{P \xrightarrow{\tau} post_{[\Delta,\tau]}(P) \;\big|\; P \subseteq Q\big\}$; this corresponds to the standard procedure that first determinizes $\mathcal{A}$ by the subset construction and then swaps its sets of final and non-final states, and $\downarrow\{Q \setminus F\}$ is the set of all subsets of $Q$ that do not contain a final state of $\mathcal{A}$. The language of $\mathcal{A}_{\mathcal{C}}$ is the complement of the language of $\mathcal{A}$, i.e. $\mathcal{L}(\mathcal{A}_{\mathcal{C}}) = \overline{\mathcal{L}(\mathcal{A})}$.

For a set of variables $\mathbb{X}$ and a variable $X$, the *projection* of $X$ from $\mathbb{X}$, denoted as $\pi_{[X]}(\mathbb{X})$, is the set $\mathbb{X} \setminus \{X\}$. For a symbol $\tau$, the projection of $X$ from $\tau$, denoted $\pi_{[X]}(\tau)$, is obtained from $\tau$ by restricting $\tau$ to the domain $\pi_{[X]}(\mathbb{X})$. For a transition relation $\Delta$, the projection of $X$ from $\Delta$, denoted as $\pi_{[X]}(\Delta)$, is the transition relation $\big\{p \xrightarrow{\pi_{[X]}(\tau)} q \mid p \xrightarrow{\tau} q \in \Delta\big\}$.

## 4 Deciding WS1S with Finite Automata

The classical decision procedure for WS1S [19] (as described in Section 3.3 of [18]) is based on a logic-automata connection and decides validity (satisfiability) of a WS1S

formula $\varphi(X_1, \ldots, X_n)$ by constructing the FA $\mathcal{A}_\varphi$ over $\{X_1, \ldots, X_n\}$ which recognizes encodings of exactly the models of $\varphi$. The automaton is built in a bottom-up manner, according to the structure of $\varphi$, starting with predefined atomic automata for literals and applying a corresponding automata operation for every logical connective and quantifier ($\land, \lor, \neg, \exists$). Hence, for every sub-formula $\psi$ of $\varphi$, the procedure will compute the automaton $\mathcal{A}_\psi$ such that $\mathcal{L}(\mathcal{A}_\psi)$ represents exactly all models of $\psi$, terminating with the result $\mathcal{A}_\varphi$.

The alphabet of $\mathcal{A}_\varphi$ consists of all symbols over the set $\mathbb{X} = \{X_1, \ldots, X_n\}$ of free variables of $\varphi$ (for $a, b \in \{0, 1\}$ and $\mathbb{X} = \{X_1, X_2\}$, we use $\begin{smallmatrix} X_1 \, : \, a \\ X_2 \, : \, b \end{smallmatrix}$ to denote the symbol $\{X_1 \mapsto a, X_2 \mapsto b\}$). A word $w$ from the language of $\mathcal{A}_\varphi$ is a sequence of these symbols, e.g. $\begin{smallmatrix} X_1 \, : \, \epsilon \\ X_2 \, : \, \epsilon \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 011 \\ X_2 \, : \, 101 \end{smallmatrix}$, or $\begin{smallmatrix} X_1 \, : \, 01100 \\ X_2 \, : \, 10100 \end{smallmatrix}$. We denote the $i$-th symbol of $w$ as $w[i]$, for $i \in \mathbb{N}_0$. An assignment $\rho : \mathbb{X} \to 2^{\mathbb{N}_0}$ mapping free variables $\mathbb{X}$ of $\varphi$ to subsets of $\mathbb{N}_0$ is encoded into a word $w_\rho$ of symbols over $\mathbb{X}$ in the following way: $w_\rho$ contains 1 in the $j$-th position of the row for $X_i$ iff $j \in X_i$ in $\rho$. Formally, for every $i \in \mathbb{N}_0$ and $X_j \in \mathbb{X}$, if $i \in \rho(X_j)$, then $w_\rho[i]$ maps $X_j \mapsto 1$. On the other hand, if $i \notin \rho(X_j)$, then either $w_\rho[i]$ maps $X_j \mapsto 0$, or the length of $w$ is smaller than or equal to $i$. Notice that there exist an infinite number of encodings of $\rho$. The shortest one is $w_\rho^s$ of the length $n + 1$, where $n$ is the largest number appearing in any of the sets that is assigned to a variable of $\mathbb{X}$ in $\rho$, or $-1$ when all these sets are empty. The rest of the encodings are all those corresponding to $w_\rho^s$ extended with an arbitrary number of $\overline{0}$ symbols appended to its end. For example, $\begin{smallmatrix} X_1 \, : \, 0 \\ X_2 \, : \, 1 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 00 \\ X_2 \, : \, 10 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 000 \\ X_2 \, : \, 100 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 000\ldots0 \\ X_2 \, : \, 100\ldots0 \end{smallmatrix}$ are all encodings of the assignment $\rho = \{X_1 \mapsto \emptyset, X_2 \mapsto \{0\}\}$. For the soundness of the decision procedure, it is important that $\mathcal{A}_\varphi$ always accepts either all encodings of $\rho$ or none of them.

The automata $\mathcal{A}_{\varphi \land \psi}$ and $\mathcal{A}_{\varphi \lor \psi}$ are constructed from $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$ by standard automata-theoretic union and intersection operations, preceded by the so-called cylindrification which unifies the alphabets of $\mathcal{A}_\varphi$ and $\mathcal{A}_\psi$. Since these operations, as well as the automata for the atomic formulae, are not the subject of the contribution proposed in this paper, we refer the interested reader to [18] for details.

The part of the procedure which is central for this paper is processing negation and existential quantification; we will therefore describe it in detail. The FA $\mathcal{A}_{\neg\varphi}$ is constructed as the complement of $\mathcal{A}_\varphi$. Then, all encodings of the assignments that were accepted by $\mathcal{A}_\varphi$ are rejected by $\mathcal{A}_{\neg\varphi}$ and vice versa. The FA $\mathcal{A}_{\exists X : \varphi}$ is obtained from the FA $\mathcal{A}_\varphi = (Q, \Delta, I, F)$ by first projecting $X$ from the transition relation $\Delta$, yielding the FA $\mathcal{A}'_\varphi = (Q, \pi_{[X]}(\Delta), I, F)$. However, $\mathcal{A}'_\varphi$ cannot be directly used as $\mathcal{A}_{\exists X : \varphi}$. The reason is that $\mathcal{A}'_\varphi$ may now be inconsistent in accepting some encodings of an assignment $\rho$ while rejecting other encodings of $\rho$. For example, suppose that $\mathcal{A}_\varphi$ accepts the words $\begin{smallmatrix} X_1 \, : \, 010 \\ X_2 \, : \, 001 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 0100 \\ X_2 \, : \, 0010 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 0100\ldots0 \\ X_2 \, : \, 0010\ldots0 \end{smallmatrix}$ and we are computing the FA for $\exists X_2 : \varphi$. When we remove $X_2$ from all symbols, we obtain $\mathcal{A}'_\varphi$ that accepts the words $\begin{smallmatrix} X_1 \, : \, 010 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 0100 \end{smallmatrix}$, $\begin{smallmatrix} X_1 \, : \, 0100\ldots0 \end{smallmatrix}$, but does not accept the word $\begin{smallmatrix} X_1 \, : \, 01 \end{smallmatrix}$ that encodes the same assignment (because $\begin{smallmatrix} X_1 \, : \, 01 \\ X_2 \, : \, ?? \end{smallmatrix} \notin \mathcal{L}(A_\varphi)$ for any values in the places of "?"s). As a remedy for this situation, we need to modify $\mathcal{A}'_\varphi$ to also accept the rest of the encodings of $\rho$. This is done by enlarging the set of final states of $\mathcal{A}'_\varphi$ to also contain all states that can reach a final state of $\mathcal{A}'_\varphi$ by a sequence of $\overline{0}$ symbols. Formally, $\mathcal{A}_{\exists X : \varphi} = (Q, \pi_{[X]}(\Delta), I, F^\sharp)$

is obtained from $\mathcal{A}'_\varphi = (Q, \pi_{[X]}(\Delta), I, F)$ by computing $F^\sharp$ from $F$ using the fixpoint computation $F^\sharp = \mu Z \,.\, F \cup pre_{[\pi_{[X]}(\Delta), \overline{0}]}(Z)$. Intuitively, the least fixpoint denotes the set of states backward-reachable from $F$ following transitions of $\pi_{[X]}(\Delta)$ labeled by $\overline{0}$.

The procedure returns an automaton $\mathcal{A}_\varphi$ that accepts exactly all encodings of the models of $\varphi$. This means that the language of $\mathcal{A}_\varphi$ is (i) universal iff $\varphi$ is valid, (ii) non-universal iff $\varphi$ is invalid, (iii) empty iff $\varphi$ is unsatisfiable, and (iv) non-empty iff $\varphi$ is satisfiable. Notice that in the particular case of *ground* formulae (i.e. formulae without free variables), the language of $\mathcal{A}_\varphi$ is either $\mathcal{L}(\mathcal{A}_\varphi) = \{\overline{0}\}^*$ in the case $\varphi$ is valid, or $\mathcal{L}(\mathcal{A}_\varphi) = \emptyset$ in the case $\varphi$ is invalid.

## 5  Nested Antichain-based Approach for Alternating Quantifiers

We now present our approach for dealing with alternating quantifiers in WS1S formulae. We consider a ground formula $\varphi$ of the form

$$\varphi = \neg\, \exists \mathcal{X}_m \,\neg \ldots \neg\, \exists \mathcal{X}_2 \underbrace{\neg\, \exists \mathcal{X}_1 : \varphi_0(\mathbb{X})}_{\varphi_1} \tag{2}$$

$$\underbrace{\phantom{\neg\, \exists \mathcal{X}_m \,\neg \ldots \neg\, \exists \mathcal{X}_2 \neg\, \exists \mathcal{X}_1 : \varphi_0(\mathbb{X})}}_{\varphi_m}$$

where each $\mathcal{X}_i$ is a set of variables $\{X_a, \ldots, X_b\}$, $\exists \mathcal{X}_i$ is an abbreviation for a non-empty sequence $\exists X_a \ldots \exists X_b$ of consecutive existential quantifications, and $\varphi_0$ is an arbitrary formula called the *matrix* of $\varphi$. Note that the problem of checking validity or satisfiability of a formula with free variables can be easily reduced to this form.

The classical procedure presented in Section 4 computes a sequence of automata $\mathcal{A}_{\varphi_0}, \mathcal{A}_{\varphi_0^\sharp}, \ldots, \mathcal{A}_{\varphi_{m-1}^\sharp}, \mathcal{A}_{\varphi_m}$ where for all $0 \le i \le m-1$, $\varphi_i^\sharp = \exists \mathcal{X}_{i+1} : \varphi_i$ and $\varphi_{i+1} = \neg \varphi_i^\sharp$. The $\varphi_i$'s are the subformulae of $\varphi$ shown in Equation (2). Since eliminating existential quantification on the automata level introduces nondeterminism (due to the projection on the transition relation), every $\mathcal{A}_{\varphi_i^\sharp}$ may be nondeterministic. The computation of $\mathcal{A}_{\varphi_{i+1}}$ then involves subset construction and becomes exponential. The worst case complexity of eliminating the prefix is therefore the tower of exponentials of the height $m$. Even though the construction may be optimized, e.g., by minimizing every $\mathcal{A}_{\varphi_i}$ (which is implemented by MONA), the size of the generated automata can quickly become intractable.

The main idea of our algorithm is inspired by the antichain algorithms [13] for testing language universality of an automaton $\mathcal{A}$. In a nutshell, testing universality of $\mathcal{A}$ is testing whether in the complement $\overline{\mathcal{A}}$ of $\mathcal{A}$ (which is created by determinization via subset construction, followed by swapping final and non-final states), an initial state can reach a final state. The crucial idea of the antichain algorithms is based on the following: (i) The search can be done on-the-fly while constructing $\overline{\mathcal{A}}$. (ii) The sets of states that arise during the search are closed (upward or downward, depending on the variant of the algorithm). (iii) The computation can be done symbolically on the generators of these closed sets. It is enough to keep only the extreme generators of the closed sets (maximal for downward, minimal for upward closed). The generators that are not extreme (we say that they are *subsumed*) can be pruned away, which vastly reduces the search space.

We notice that individual steps of the algorithm for constructing $\mathcal{A}_\varphi$ are very similar to testing universality. Automaton $\mathcal{A}_{\varphi_i}$ arises by subset construction from $\mathcal{A}_{\varphi_{i-1}^\sharp}$, and to compute $\mathcal{A}_{\varphi_i^\sharp}$, it is necessary to compute the set of final states $F_i^\sharp$. Those are states backward reachable from the final states of $\mathcal{A}_{\varphi_i}$ via a subset of transitions of $\Delta_i$ (those labelled by symbols projected to $\overline{0}$ by $\pi_{i+1}$). To compute $F_i^\sharp$, the antichain algorithms could be actually taken off-the-shelf and run with $\mathcal{A}_{\varphi_{i-1}^\sharp}$ in the role of the input $\mathcal{A}$ and $\mathcal{A}_{\varphi_i^\sharp}$ in the role of $\overline{\mathcal{A}}$. However, this approach has the following two problems. First, antichain algorithms do not produce the automaton $\overline{\mathcal{A}}$ (here $\mathcal{A}_{\varphi_i^\sharp}$), but only a symbolic representation of a set of (backward) reachable states (here of $F_i^\sharp$). Since $\mathcal{A}_{\varphi_i^\sharp}$ is the input of the construction of $\mathcal{A}_{\varphi_{i+1}}$, the construction of $\mathcal{A}_\varphi$ could not continue. The other problem is that the size of the input $\mathcal{A}_{\varphi_{i-1}^\sharp}$ of the antichain algorithm is only limited by the tower of exponentials of the height $i-1$, and this might be already far out of reach.

The main contribution of our paper is an algorithm that alleviates the two problems mentioned above. It is based on a novel way of performing not only one, but all the $2m$ steps of the construction of $\mathcal{A}_\varphi$ on-the-fly. It uses a nested symbolic representation of sets of states and a form of nested subsumption pruning on all levels of their structure. This is achieved by a substantial refinement of the basic ideas of antichain algorithms.

### 5.1 Structure of the Algorithm

Let us now start explaining our on-the-fly algorithm for handling quantifier alternation. Following the construction of automata in Section 4, the structure of the automata from the previous section, $\mathcal{A}_{\varphi_0}, \mathcal{A}_{\varphi_0^\sharp}, \ldots, \mathcal{A}_{\varphi_{m-1}^\sharp}, \mathcal{A}_{\varphi_m}$, can be described using the following recursive definition. We use $\pi_i(C)$ for any mathematical structure $C$ to denote projection of all variables in $\mathcal{X}_1 \cup \cdots \cup \mathcal{X}_i$ from $C$.

Let $\mathcal{A}_{\varphi_0} = (Q_0, \Delta_0, I_0, F_0)$ be an FA over $\mathbb{X}$. Then, for each $0 \le i < m$, $\mathcal{A}_{\varphi_i^\sharp}$ and $\mathcal{A}_{\varphi_{i+1}}$ are FAs over $\pi_{i+1}(\mathbb{X})$ that have from the construction the following structure:

$$
\begin{array}{l|l}
\mathcal{A}_{\varphi_i^\sharp} = (Q_i, \Delta_i^\sharp, I_i, F_i^\sharp) \text{ where} & \mathcal{A}_{\varphi_{i+1}} = (Q_{i+1}, \Delta_{i+1}, I_{i+1}, F_{i+1}) \text{ where} \\
\quad \Delta_i^\sharp = \pi_{i+1}(\Delta_i) \text{ and} & \quad \Delta_{i+1} = \left\{ R \xrightarrow{\tau} post[\Delta_i^\sharp, \tau](R) \,\middle|\, R \in Q_{i+1} \right\}, \\
\quad F_i^\sharp = \mu Z \,.\, F_i \cup pre[\Delta_i^\sharp, \overline{0}](Z). & \quad Q_{i+1} = 2^{Q_i}, \quad I_{i+1} = \{I_i\}, \quad \text{and} \quad F_{i+1} = \Uparrow\{Q_i \setminus F_i^\sharp\}.
\end{array}
$$

We recall that $\mathcal{A}_{\varphi_i^\sharp}$ directly corresponds to existential quantification (cf. Section 4), and $\mathcal{A}_{\varphi_{i+1}}$ directly corresponds to the complement of $\mathcal{A}_{\varphi_i^\sharp}$ (cf. Section 3).

A crucial observation behind our approach is that, because $\varphi$ is ground, $\mathcal{A}_\varphi$ is an FA over an empty set of variables, and, therefore, $\mathcal{L}(\mathcal{A}_\varphi)$ is either the empty set $\emptyset$ or the set $\{\overline{0}\}^*$. Therefore, we need to distinguish between these two cases only. To determine which of them holds, we do not need to explicitly construct the automaton $\mathcal{A}_\varphi$. Instead, it suffices to check whether $\mathcal{A}_\varphi$ accepts the empty string $\epsilon$. This is equivalent to checking existence of a state that is at the same time final and initial, that is

$$\models \varphi \quad \text{iff} \quad I_m \cap F_m \neq \emptyset. \tag{3}$$

To compute $I_m$ from $I_0$ is straightforward (it equals $\{\{\ldots\{\{I_0\}\}\ldots\}\}$ nested $m$-times). In the rest of the section, we will describe how to compute $F_m$ (its symbolic representation), and how to test whether it intersects with $I_m$.

7

The algorithm takes advantage of the fact that to represent final states, one can use their complement, the set of non-final states. For $0 \leq i \leq m$, we write $N_i$ and $N_i^\sharp$ to denote the sets of non-final states $Q_i \setminus F_i$ of $\mathcal{A}_i$ and $Q_i \setminus F_i^\sharp$ of $\mathcal{A}_i^\sharp$, respectively. The algorithm will then instead of computing the sequence of automata $\mathcal{A}_{\varphi_0}, \mathcal{A}_{\varphi_0^\sharp}, \ldots,$ $\mathcal{A}_{\varphi_{m-1}^\sharp}, \mathcal{A}_{\varphi_m}$ compute the sequence $F_0, F_0^\sharp, N_1, N_1^\sharp, \ldots$ up to either $F_m$ (if $m$ is even) or $N_m$ (if $m$ is odd), which suffices for testing the validity of $\varphi$. The algorithm starts with $F_0$ and uses the following recursive equations:

$$\text{(i) } F_{i+1} = \downarrow\{N_i^\sharp\} \qquad \text{(ii) } F_i^\sharp = \mu Z \,.\, F_i \cup pre_{[\Delta_i^\sharp, \overline{0}]}(Z)$$
$$\text{(iii) } N_{i+1} = \uparrow\textstyle\coprod\{F_i^\sharp\} \quad \text{(iv) } N_i^\sharp = \nu Z \,.\, N_i \cap cpre_{[\Delta_i^\sharp, \overline{0}]}(Z) \tag{4}$$

Intuitively, Equations (i) and (ii) are directly from the definition of $\mathcal{A}_i$ and $\mathcal{A}_i^\sharp$. Equation (iii) is a dual of Equation (i): $N_{i+1}$ contains all subsets of $Q_i$ that contain at least one state from $F_i^\sharp$. Finally, Equation (iv) is a dual of Equation (ii): in the $k$-th iteration of the greatest fixpoint computation, the current set of states $Z$ will contain all states that cannot reach an $F_i$ state over $\overline{0}$ within $k$ steps. In the next iteration, only those states of $Z$ are kept such that all their $\overline{0}$-successors are in $Z$. Hence, the new value of $Z$ is the set of states that cannot reach $F_i$ over $\overline{0}$ in $k+1$ steps, and the computation stabilises with the set of states that cannot reach $F_i$ over $\overline{0}$ in any number of steps.

In the next two sections, we will show that both of the above fixpoint computations can be carried out symbolically on representatives of upward/downward closed sets. Particularly, in Sections 5.2 and 5.3, we show how the fixpoints from Equations (ii) and (iv) can be computed symbolically, using subsets of $Q_{i-1}$ as representatives (generators) of upward/downward closed subsets of $Q_i$. Section 5.4 explains how the above symbolic fixpoint computations can be carried out using nested terms of depth $i$ as a symbolic representation of computed states of $Q_i$. Section 5.5 shows how to test emptiness of $I_m \cap F_m$ on the symbolic terms, and Section 5.6 describes the subsumption relation used to minimize the symbolic term representation used within computations of Equations (ii) and (iv). Proofs of the lemmas and used equations can be found in [25].

## 5.2  Computing $N_i^\sharp$ on Representatives of $\uparrow\coprod\mathcal{R}$-sets

Computing $N_i^\sharp$ at each odd level of the hierarchy of automata is done by computing the greatest fixpoint of the function from Equation (4iv):

$$f_{N_i^\sharp}(Z) = N_i \cap cpre_{[\Delta_i^\sharp, \overline{0}]}(Z). \tag{5}$$

We will show that the whole fixpoint computation from Equation (4iv) can be carried out symbolically on the representatives of $Z$. We will explain that: (a) All intermediate values of $Z$ have the form $\uparrow\coprod\mathcal{R}$, $\mathcal{R} \subseteq Q_i$, so the sets $\mathcal{R}$ can be used as their symbolic representatives. (b) $cpre$ and $\cap$ can be computed on such a representation efficiently.

Let us start with the computation of $cpre_{[\Delta_i^\sharp, \tau]}(Z)$ where $\tau \in \pi_{i+1}(\mathbb{X})$, assuming that $Z$ is of the form $\uparrow\coprod\mathcal{R}$, represented by $\mathcal{R} = \{R_1, \ldots, R_n\}$. Observe that a set of symbolic representatives $\mathcal{R}$ stands for the intersection of denotations of individual representatives, that is

$$\uparrow\textstyle\coprod\mathcal{R} = \bigcap_{R_j \in \mathcal{R}} \uparrow\textstyle\coprod\{R_j\}. \tag{6}$$

$Z$ can thus be written as the *cpre*-image $cpre[\Delta_i^\sharp,\tau](\bigcap \mathcal{S})$ of the intersection of the elements of a set $\mathcal{S}$ having the form $\uparrow\coprod\{R_j\}$, $R_j \in \mathcal{R}$. Further, because *cpre* distributes over $\cap$, we can compute the *cpre*-image of an intersection by computing intersection of the *cpre*-images, i.e.,

$$cpre[\Delta_i^\sharp,\tau](\bigcap \mathcal{S}) = \bigcap_{S \in \mathcal{S}} cpre[\Delta_i^\sharp,\tau](S). \tag{7}$$

By the definition of $\Delta_i^\sharp$ (where $\Delta_i^\sharp = \pi_{i+1}(\Delta_i)$), $cpre[\Delta_i^\sharp,\tau](S)$ can be computed using the transition relation $\Delta_i$ for the price of further refining the intersection. In particular,

$$cpre[\Delta_i^\sharp,\tau](S) = \bigcap_{\omega \in \pi_{i+1}^{-1}(\tau)} cpre[\Delta_i,\omega](S). \tag{8}$$

Intuitively, $cpre[\Delta_i^\sharp,\tau](S)$ contains states from which every transition labelled by *any* symbol that is projected to $\tau$ by $\pi_{i+1}$ has its target in $S$. Using Equations (6), (7), and (8), we can write $cpre[\Delta_i^\sharp,\tau](Z)$ as $\bigcap_{S \in \mathcal{S},\omega \in \pi_{i+1}^{-1}(\tau)} cpre[\Delta_i,\omega](S)$.

To compute the individual conjuncts $cpre[\Delta_i,\omega](S)$, we take advantage of the fact that every $S$ is in the special form $\uparrow\coprod\{R_j\}$, and that $\Delta_i$ is, by its definition (determinization via subset construction), *monotone* w.r.t. $\subseteq$. That is, if $P \xrightarrow{\omega} P' \in \Delta_i$ for some $P, P' \in Q_i$, then for every $R \supseteq P$, there is $R' \supseteq P'$ s.t. $R \xrightarrow{\omega} R' \in \Delta_i$. Due to monotonicity, the $cpre[\Delta_i,\omega]$-image of an upward closed set is also upward closed. Moreover, we observe that it can be computed symbolically using *pre* on elements of its generators. Particularly, for a set of singletons $S = \uparrow\coprod\{R_j\}$, we get the following equation:

$$cpre[\Delta_i,\omega](\uparrow\coprod\{R_j\}) = \uparrow\coprod\left\{pre[\Delta_{i-1}^\sharp,\omega](R_j)\right\}. \tag{9}$$

Intuitively, the sets with *post*-images above a singleton $\{p\} \in \{\{p\} \mid p \in R_j\} = \uparrow\coprod\{R_j\}$ are those that contain at least one state $q \in Q_{i-1}$ s.t. $q \xrightarrow{\omega} p \in \Delta_{i-1}^\sharp$. Using (9), $cpre[\Delta_i^\sharp,\tau](Z)$ can be rewritten as $\bigcap_{R \in \mathcal{R},\omega \in \pi_{i+1}^{-1}(\tau)} \uparrow\coprod\left\{pre[\Delta_{i-1}^\sharp,\omega](R_j)\right\}$. By applying (6), we get the final formula for $cpre[\Delta_i^\sharp,\tau]$ shown in the lemma below.

**Lemma 1.** $cpre[\Delta_i^\sharp,\tau](\uparrow\coprod\mathcal{R}) = \uparrow\coprod\left\{pre[\Delta_{i-1}^\sharp,\omega](R_j) \mid \omega \in \pi_{i+1}^{-1}(\tau), R_j \in \mathcal{R}\right\}$.

To compute $f_{N_i^\sharp}(Z)$, it remains to intersect $cpre[\Delta_i^\sharp,\overline{0}](Z)$, computed using Lemma 1, with $N_i$. By Equation (4iii), $N_i$ equals $\uparrow\coprod\{F_{i-1}^\sharp\}$, and, by Equation (6), the intersection can be done symbolically as

$$f_{N_i^\sharp}(Z) = \uparrow\coprod\left(\{F_{i-1}^\sharp\} \cup \left\{pre[\Delta_{i-1}^\sharp,\omega](R_j) \mid \omega \in \pi_{i+1}^{-1}(\overline{0}), R_j \in \mathcal{R}\right\}\right). \tag{10}$$

Finally, note that a symbolic application of $f_{N_i^\sharp}$ to $Z = \uparrow\coprod\mathcal{R}$ represented as the set $\mathcal{R}$ reduces to computing *pre*-images of the elements of $\mathcal{R}$, which are then put next to each other, together with $F_{i-1}^\sharp$. The computation starts from $N_i = \uparrow\coprod\{F_{i-1}^\sharp\}$, represented by $\{F_{i-1}^\sharp\}$, and each of its steps, implemented by Equation (10), preserves the form of sets $\uparrow\coprod\mathcal{R}$, represented by $\mathcal{R}$.

### 5.3 Computing $F_i^\sharp$ on Representatives of $\downarrow \mathcal{R}$-sets

Similarly as in the previous section, computation of $F_i^\sharp$ at each even level of the automata hierarchy is done by computing the least fixpoint of the function

$$f_{F_i^\sharp}(Z) = F_i \cup pre[\Delta_i^\sharp, \overline{0}](Z). \tag{11}$$

We will show that the whole fixpoint computation from Equation (4ii) can be carried out symbolically. We will explain the following: (a) All intermediate values of $Z$ are of the form $\downarrow \mathcal{R}$, $\mathcal{R} \subseteq Q_i$, so the sets $\mathcal{R}$ can be used as their symbolic representatives. (b) $pre$ and $\cup$ can be computed efficiently on such a symbolic representation. The computation is a simpler analogy of the one in Section 5.2.

We start with the computation of $pre[\Delta_i^\sharp, \tau](Z)$ where $\tau \in \pi_{i+1}(\mathbb{X})$, assuming that $Z$ is of the form $\downarrow \mathcal{R}$, represented by $\mathcal{R} = \{R_1, \dots, R_n\}$. A simple analogy to Equations (6) and (7) of Section 5.2 is that the union of downward closed sets is a downward closed set generated by the union of their generators, i.e., $\downarrow \mathcal{R} = \bigcup_{R_j \in \mathcal{R}} \downarrow\{R_j\}$ and that $pre$ distributes over union, i.e., $pre[\Delta_i^\sharp, \tau](\bigcup \mathcal{R}) = \bigcup_{R_j \in \mathcal{R}} pre[\Delta_i^\sharp, \tau](\downarrow\{R_j\})$. An analogy of Equation (8) holds too:

$$pre[\Delta_i^\sharp, \tau](S) = \bigcup_{\omega \in \pi_{i+1}^{-1}(\tau)} pre[\Delta_i, \omega](S). \tag{12}$$

Intuitively, $pre[\Delta_i^\sharp, \tau](S)$ contains states from which *at least one* transition labelled by *any* symbol that is projected to $\tau$ by $\pi_{i+1}$ leaves with the target in $S$. Using Equation (12), we can write $pre[\Delta_i^\sharp, \tau](Z)$ as $\bigcup_{R_j \in \mathcal{R}, \omega \in \pi_{i+1}^{-1}(\tau)} pre[\Delta_i, \omega](\downarrow\{R_j\})$.

To compute the individual disjuncts $pre[\Delta_i, \omega](\downarrow\{R_j\})$, we take advantage of the fact that every $\downarrow\{R_j\}$ is downward closed, and that $\Delta_i$ is, by definition (determinization by subset construction), *monotone* w.r.t. $\supseteq$. That is, if $P \xrightarrow{\omega} P' \in \Delta_i$ for some $P, P' \in Q_i$, then for every $R \subseteq P$, there is $R' \subseteq P'$ s.t. $R \xrightarrow{\omega} R' \in \Delta_i$. Due to monotonicity, the $pre[\Delta_i, \omega]$-image of a downward closed set is downward closed. Moreover, we observe that it can be computed symbolically using $cpre$ on elements of its generators. In particular, for a set $\downarrow\{R_j\}$, we get the following equation, which is a dual of Equation (9):

$$pre[\Delta_i, \omega](\downarrow\{R_j\}) = \downarrow\{cpre[\Delta_{i-1}^\sharp, \omega](R_j)\}. \tag{13}$$

Intuitively, the sets with the *post*-images below $R_j$ are those which do not have an outgoing transition leading outside $R_j$. The largest such set is $cpre[\Delta_{i-1}^\sharp, \omega](R_j)$. Using Equation (13), $pre[\Delta_i^\sharp, \tau](Z)$ can be rewritten as $\bigcup_{R_j \in \mathcal{R}, \omega \in \pi_{i+1}^{-1}(\tau)} \downarrow\{cpre[\Delta_{i-1}^\sharp, \omega](R_j)\}$, which gives us the final formula for $pre[\Delta_i^\sharp, \tau]$ described in Lemma 2.

**Lemma 2.** $pre[\Delta_i^\sharp, \tau](\downarrow \mathcal{R}) = \downarrow\{cpre[\Delta_{i-1}^\sharp, \omega](R_j) \mid \omega \in \pi_{i+1}^{-1}(\tau), R_j \in \mathcal{R}\}$.

To compute $f_{F_i^\sharp}(Z)$, it remains to unite $pre[\Delta_i^\sharp, \overline{0}](Z)$, computed using Lemma 2, with $F_i$. From Equation (4i), $F_i$ equals $\downarrow\{N_{i-1}^\sharp\}$, so the union can be done symbolically as

$$f_{F_i^\sharp}(Z) = \downarrow\left(\{N_{i-1}^\sharp\} \cup \{cpre[\Delta_{i-1}^\sharp, \omega](R_j) \mid \omega \in \pi_{i+1}^{-1}(\overline{0}), R_j \in \mathcal{R}\}\right). \tag{14}$$

Therefore, a symbolic application of $f_{F_i^\sharp}$ to $Z = \downarrow \mathcal{R}$ represented using the set $\mathcal{R}$ reduces to computing $cpre$-images of elements of $\mathcal{R}$, which are put next to each other,

together with $N_{i-1}^{\sharp}$. The computation starts from $F_i = \downarrow\{N_{i-1}^{\sharp}\}$, represented by $N_{i-1}^{\sharp}$, and each of its steps, implemented by Equation (14), preserves the form of sets $\downarrow\mathcal{R}$, represented by $\mathcal{R}$.

## 5.4 Computation of $F_i^{\sharp}$ and $N_i^{\sharp}$ on Symbolic Terms

Sections 5.2 and 5.3 show how sets of states arising within the fixpoint computations from Equations (4ii) and (4iv) can be represented symbolically using representatives which are sets of states of the lower level. The sets of states of the lower level will be again represented symbolically. When computing the fixpoint of level $i$, we will work with nested symbolic representation of states of depth $i$. Particularly, sets of states of $Q_k, 0 \leq k \leq i$, are represented by *terms of level $k$* where a term of level 0 is a subset of $Q_0$, a term of level $2j + 1$, $j \geq 0$, is of the form $\uparrow\coprod\{t_1, \ldots, t_n\}$ where $t_1, \ldots, t_n$ are terms of level $2j$, and a term of level $2j$, $j > 0$, is of the form $\downarrow\{t_1, \ldots, t_n\}$ where $t_1, \ldots, t_n$ are terms of level $2j - 1$.

The computation of *cpre* and $f_{N_{2j+1}^{\sharp}}$ on a term of level $2j + 1$ and computation of *pre* and $f_{F_{2j}^{\sharp}}$ on a term of level $2j$ then becomes a recursive procedure that descends via the structure of the terms and produces again a term of level $2j + 1$ or $2j$, respectively. In the case of *cpre* and $f_{N_{2j+1}^{\sharp}}$ called on a term of level $2j + 1$, Lemma 1 reduces the computation to a computation of *pre* on its sub-terms of level $2j$, which is again reduced by Lemma 2 to a computation of *cpre* on terms of level $2j - 1$, and so on until the bottom level where the algorithm computes *pre* on the terms of level 0 (subsets of $Q_0$). The case of *pre* and $f_{F_{2j}^{\sharp}}$ called on a term of level $2j$ is symmetrical.

*Example* We will demonstrate the run of our algorithm on the following abstract example. Consider a ground WS1S formula $\varphi = \neg\exists\mathcal{X}_3\neg\exists\mathcal{X}_2\neg\exists\mathcal{X}_1 : \varphi_0$ and an FA $\mathcal{A}_0 = (Q_0, \Delta_0, I_0 = \{a\}, F_0 = \{a, b\})$ that represents $\varphi_0$. Recall that our method decides validity of $\varphi$ by computing symbolically the sequence of sets $F_0^{\sharp}, N_1, N_1^{\sharp}, F_2, F_2^{\sharp}, N_3$, each of them represented using a symbolic term, and then checks if $I_3 \cap N_3 \neq \emptyset$. In the following paragraph, we will show how such a sequence is computed and interleave the description with examples of possible intermediate results.

The fixpoint computation from Equation (4ii) of the first set in the sequence, $F_0^{\sharp}$, is an explicit computation of the set of states backward-reachable from $F_0$ via $\overline{0}$ transitions of $\Delta_0^{\sharp}$. It is done using Equation (12), yielding, e.g., the term

$$t_{[F_0^{\sharp}]} = F_0^{\sharp} = \{a, b, c\}.$$

The fixpoint computation of $N_1^{\sharp}$ from Equation (4iv) of is done symbolically. It starts from $N_1$ represented using Equation (4iii) as the term $t_{[N_1]} = \uparrow\coprod\{\{a, b, c\}\}$, and each of its iterations is carried out using Equation (10). Equation (10) transforms the problem of computing $cpre_{[\Delta_1, \omega']}$-image of a term into a computation of a series of $pre_{[\Delta_0^{\sharp}, \omega]}$-images of its sub-terms, which is carried out using Equation (12) in the same way as when computing $t_{[F_0^{\sharp}]}$, ending with, e.g., the term

$$t_{[N_1^{\sharp}]} = \uparrow\coprod\{\{a, b, c\}, \{b, c\}, \{c, d\}\}.$$

The term representing $F_2$ is then $t_{[F_2]} = \downarrow\{t_{[N_1^{\sharp}]}\}$, due to Equation (4i). The symbolic fixpoint computation of $F_2^{\sharp}$ from Equation (4ii) then starts from $t_{[F_2]}$, in our example

$$t_{[F_2]} = \downarrow\left\{\uparrow\coprod\{\{a, b, c\}, \{b, c\}, \{c, d\}\}\right\}.$$

Its steps are computed using Equation (14), which transforms the computation of the image of $pre_{[\Delta_2^\sharp,\omega'']}$ into computations of a series of $cpre_{[\Delta_1^\sharp,\omega']}$-images of sub-terms. These are in turn transformed by Lemma 1 into computations of $pre_{[\Delta_0^\sharp,\omega]}$-images of sub-sub-terms, subsets of $Q_0$, in our example yielding, e.g., the term

$$t_{[F_2^\sharp]} = \downarrow\Big\{\uparrow\coprod\{\{a,b,c\},\{b,c\},\{c,d\}\},\uparrow\coprod\{\{b\},\{d\}\},\uparrow\coprod\{\{a\},\{c,d\}\}\Big\}.$$

Using Equation (4iv), the final term representing $N_3$ is then

$$t_{[N_3]} = \uparrow\coprod\Big\{\downarrow\big\{\uparrow\coprod\{\{a,b,c\},\{b,c\},\{c,d\}\},\uparrow\coprod\{\{b\},\{d\}\},\uparrow\coprod\{\{a\},\{c,d\}\}\big\}\Big\}.$$

In the next section, we will describe how we check whether $I_3 \cap F_3 \neq \emptyset$ using the computed term $t_{[N_3]}$.

## 5.5 Testing $I_m \cap F_m \overset{?}{\neq} \emptyset$ on Symbolic Terms

Due to the special form of the set $I_m$ (every $I_i, 1 \leq i \leq m$, is the singleton $\{I_{i-1}\}$), the test $I_m \cap F_m \neq \emptyset$ can be done efficiently over the symbolic terms representing $F_m$. Because $I_m = \{I_{m-1}\}$ is a singleton set, testing $I_m \cap F_m \neq \emptyset$ is equivalent to testing $I_{m-1} \in F_m$. If $m$ is odd, our approach computes the symbolic representation of $N_m$ instead of $F_m$. Obviously, since $N_m$ is the complement of $F_m$, it holds that $I_{m-1} \in F_m \iff I_{m-1} \notin N_m$. Our way of testing $I_{m-1} \in Y_m$ on a symbolic representation of the set $Y_m$ of level $m$ is based on the following equations:

$$\{x\} \in \downarrow\mathbb{Y} \iff \exists Y \in \mathbb{Y} : x \in Y \tag{15}$$

$$\{x\} \in \uparrow\coprod\mathbb{Y} \iff \forall Y \in \mathbb{Y} : x \in Y \tag{16}$$

and for $i = 0$, $\quad\quad I_0 \in \uparrow\coprod\mathbb{Y} \iff \forall Y \in \mathbb{Y} : I_0 \cap Y \neq \emptyset. \tag{17}$

Given a symbolic term $t_{[X]}$ of level $m$ representing a set $X \subseteq Q_m$, testing emptiness of $I_m \cap F_m$ or $I_m \cap N_m$ can be done over $t_{[X]}$ by a recursive procedure that descends along the structure of $t_{[X]}$ using Equations (15) and (16), essentially generating an AND-OR tree, terminating the descent by the use of Equation (17).

*Example* In the example of Section 5.4, we would test whether $\{\{\{\{a\}\}\}\} \cap N_3 = \emptyset$ over $t_{[N_3]}$. This is equivalent to testing whether $I_2 = \{\{\{a\}\}\} \in N_3$. From Equation (16) we get that $I_2 \in N_3 \iff I_1 = \{\{a\}\} \in F_2^\sharp$ because $F_2^\sharp$ is the denotation of the only sub-term $t_{[F_2^\sharp]}$ of $t_{[N_3]}$. Equation (15) implies that $I_1 = \{\{a\}\} \in F_2^\sharp \iff \{a\} \in N_1^\sharp \vee \{a\} \in \uparrow\coprod\{\{b\},\{d\}\} \vee \{a\} \in \uparrow\coprod\{\{a\},\{c,d\}\}$. Each of the disjuncts could then be further reduced by Equation (16) into a conjunction of membership queries on the base level which would be solved by Equation (17). Since none of the disjuncts is satisfied, we conclude that $I_1 \notin F_2^\sharp$, so $I_2 \notin N_3$, implying that $I_2 \in F_3$ and thus obtain the result $\models \varphi$.

## 5.6 Subsumption of Symbolic Terms

Although the use of symbolic terms instead of an explicit enumeration of sets of states itself considerably reduces the searched space, an even greater degree of reduction can be obtained using subsumption inside the symbolic representatives to reduce their size, similarly as in the antichain algorithms [14]. For any set of sets $\mathbb{X}$ containing a pair of distinct elements $Y, Z \in \mathbb{X}$ s.t. $Y \subseteq Z$, it holds that

$$\downarrow\mathbb{X} = \downarrow(\mathbb{X} \setminus Y) \quad \text{and} \quad \uparrow\coprod\mathbb{X} = \uparrow\coprod(\mathbb{X} \setminus Z). \tag{18}$$

Therefore, if $\mathbb{X}$ is used to represent the set $\downarrow\mathbb{X}$, the element $Y$ is *subsumed* by $Z$ and can be removed from $\mathbb{X}$ without changing its denotation. Likewise, if $\mathbb{X}$ is used to represent $\uparrow\coprod\mathbb{X}$, the element $Z$ is *subsumed* by $Y$ and can be removed from $\mathbb{X}$ without changing its denotation. We can thus simplify any symbolic term by pruning out its sub-terms that represent elements subsumed by elements represented by other sub-terms, without changing the denotation of the term.

Computing subsumption on terms can be done using the following two equations:

$$\downarrow\mathbb{X} \subseteq \downarrow\mathbb{Y} \qquad \Longleftrightarrow \qquad \forall X \in \mathbb{X} \exists Y \in \mathbb{Y} : X \subseteq Y \qquad (19)$$

$$\uparrow\coprod\mathbb{X} \subseteq \uparrow\coprod\mathbb{Y} \qquad \Longleftrightarrow \qquad \forall Y \in \mathbb{Y} \exists X \in \mathbb{X} : X \subseteq Y. \qquad (20)$$

Using Equations (19) and (20), testing subsumption of terms of level $i$ reduces to testing subsumption of terms of level $i-1$. The procedure for testing subsumption of two terms descends along the structure of the term, using (19) and (20) on levels greater than 0, and on level 0, where terms are subsets of $Q_0$, it tests subsumption by set inclusion.

*Example* In the example from Section 5.4, we can use the inclusion $\{b, c\} \subseteq \{a, b, c\}$ and Equation (18) to reduce $t_{[N_1^\sharp]} = \uparrow\coprod\{\{a, b, c\}, \{b, c\}, \{c, d\}\}$ to the term

$$t_{[N_1]}' = \uparrow\coprod\{\{b, c\}, \{c, d\}\}.$$

Moreover, Equation (20) implies that $\uparrow\coprod\{\{b, c\}, \{c, d\}\}$ is subsumed by the term $\uparrow\coprod\{\{b\}, \{d\}\}$, and, therefore, we can reduce the term $t_{[F_2^\sharp]}$ to the term

$$t_{[F_2^\sharp]}' = \downarrow\left\{\uparrow\coprod\{\{b\}, \{d\}\}, \uparrow\coprod\{\{a\}, \{c, d\}\}\right\}.$$

## 6 Experimental Evaluation

We implemented a prototype of the presented approach in the tool `dWiNA` [20] and evaluated it in a benchmark of both practical and generated examples. The tool uses the frontend of MONA to parse input formulae and also for the construction of the base automaton $\mathcal{A}_{\varphi_0}$, and further uses the MTBDD-based representation of FAs from the `libvata` [21] library. The tool supports the following two modes of operation.

In mode I, we use MONA to generate the deterministic automaton $\mathcal{A}_{\varphi_0}$ corresponding to the matrix of the formula $\varphi$, translate it to `libvata` and run our algorithm for handling the prefix of $\varphi$ using `libvata`. In mode II, we first translate the formula $\varphi$ into the formula $\varphi'$ in prenex normal form (i.e., it consists of a quantifier prefix and a quantifier-free matrix) where the occurence of negation in the matrix is limited to literals, and then construct the nondeterministic automaton $\mathcal{A}_{\varphi_0}$ directly using `libvata`.

Our experiments were performed on an Intel Core i7-4770@3.4 GHz processor with 32 GiB RAM. The practical formulae for our experiments that we report on here were obtained from the shape analysis of [5] and evaluated using mode I of our tool; the results are shown in Table 1 (see [20] for additional experimental results). We measure the time of runs of the tools for processing only the prefix of the formulae. We can observe that w.r.t. the speed, we get comparable results; in some cases `dWiNA` is slower than MONA, which we attribute to the fact that our prototype implementation is, when compared with MONA, quite immature. Regarding space, we compare the sum of the number of states of all automata generated by MONA when processing the prefix of $\varphi$ with the number of symbolic terms generated by `dWiNA` for processing the same. We can observe a significant reduction in the generated state space.

We also tried to run `dWiNA` on the modified formulae in mode II but ran into the problem that we were not able to construct the non-deterministic automaton for the quantifier-free matrix $\varphi_0$. This was because after transformation of $\varphi$ into prenex

**Table 1.** Results for practical examples

| Benchmark | Time [s] | | Space [states] | |
|---|---|---|---|---|
| | MONA | dWiNA | MONA | dWiNA |
| reverse-before-loop | 0.01 | 0.01 | 179 | 47 |
| insert-in-loop | 0.01 | 0.01 | 463 | 110 |
| bubblesort-else | 0.01 | 0.01 | 1 285 | 271 |
| reverse-in-loop | 0.02 | 0.02 | 1 311 | 274 |
| bubblesort-if-else | 0.02 | 0.23 | 4 260 | 1 040 |
| bubblesort-if-if | 0.12 | 1.14 | 8 390 | 2 065 |

normal form, if $\varphi_0$ contains many conjunctions, the sizes of the automata generated using intersection grow too large (one of the reasons for this is that `libvata` in its current version does not support efficient reduction of automata).

To better evaluate the scalability of our approach, we created several parameterized families of WS1S formulae. We start with basic formulae encoding interesting relations among subsets of $\mathbb{N}_0$, such as existence of certain transitive relations, singleton sets, or intervals (their full definition can be found in [20]). From these we algorithmically create families of formulae with

**Table 2.** Results for generated formulae

| $k$ | Time [s] | | Space [states] | |
|---|---|---|---|---|
| | MONA | dWiNA | MONA | dWiNA |
| 2 | 0.20 | 0.01 | 25 517 | 44 |
| 3 | 0.57 | 0.01 | 60 924 | 50 |
| 4 | 1.79 | 0.02 | 145 765 | 58 |
| 5 | 4.98 | 0.02 | 349 314 | 70 |
| 6 | $\infty$ | 0.47 | $\infty$ | 90 |

larger quantifier depth, regardless of the meaning of the created formulae (though their semantics is still nontrivial). In Table 2, we give the results for one of the families where the basic formula expresses existence of an ascending chain of $n$ sets ordered w.r.t. $\subset$. The parameter $k$ stands for the number of alternations in the prefix of the formulae:

$$\exists Y : \neg \exists X_1 \neg \ldots \neg \exists X_k, \ldots, X_n : \bigwedge_{1 \leq i < n} \left( X_i \subseteq Y \wedge X_i \subset X_{i+1} \right) \Rightarrow X_{i+1} \subseteq Y.$$

We ran the experiments in mode II of `dWiNA` (the experiment in mode I was not successful due to a costly conversion of a large base automaton from MONA to `libvata`).

## 7 Conclusion and Future Work

We presented a new approach for dealing with alternating quantifications within the automata-based decision procedure for WS1S. Our approach is based on a generalization of the idea of the so-called antichain algorithm for testing universality or language inclusion of finite automata. Our approach processes a prefix of the formula with an arbitrary number of quantifier alternations on-the-fly using an efficient symbolic representation of the state space, enhanced with subsumption pruning. Our experimental results are encouraging (our tool often outperforms MONA) and show that the direction started in this paper—using modern techniques for nondeterministic automata in the context of deciding WS1S formulae—is promising.

An interesting direction of further development seems to be lifting the symbolic *pre*/*cpre* operators to a more general notion of terms that allow working with general sub-formulae (that may include logical connectives and nested quantifiers). The algorithm could then be run over arbitrary formulae, without the need of the transformation into the prenex form. This would open a way of adopting optimizations used in other tools as well as syntactical optimizations of the input formula such as anti-prenexing.

Another way of improvement is using simulation-based techniques to reduce the generated automata as well as to weaken the term-subsumption relation (an efficient algorithm for computing simulation over BDD-represented automata is needed). We also plan to extend the algorithms to WS$k$S and tree-automata, perhaps even further to more general inductive structures.

## References

1. Meyer, A.R.: Weak monadic second order theory of successor is not elementary-recursive. *In Proc. of Logic Colloquium—Symposium on Logic Held at Boston*. LNCS 453, Springer, 1972.
2. Elgaard, J., Klarlund, N., Møller, A.: MONA 1.x: new techniques for WS1S and WS2S. *In Proc. of CAV'98*. LNCS 1427, Springer, 1998.
3. Klarlund, N., Møller, A.: MONA Ver. 1.4 Manual. Available from http://www.brics.dk/mona/.
4. Madhusudan, P., Parlato, G., Qiu, X.: Decidable logics combining heap structures and data. *In Proc. of POPL'11*. ACM, 2011.
5. Madhusudan, P., Qiu, X.: Efficient decision procedures for heaps using STRAND. *In Proc. of SAS'11*. LNCS 6887, Springer, 2011.
6. Iosif, R., Rogalewicz, A., Šimáček, J.: The tree width of separation logic with recursive definitions. *In Proc. of CADE'13*. LNCS 7898, Springer, 2013.
7. Chin, W., David, C., Nguyen, H.H., Qin, S.: Automated verification of shape, size and bag properties via user-defined predicates in separation logic. Science of Computer Programing **77**(9), 2012.
8. Zee, K., Kuncak, V., Rinard, M.C.: Full functional verification of linked data structures. *In Proc. of POPL'08*. ACM, 2008.
9. Hamza, J., Jobstmann, B., Kuncak, V.: Synthesis for regular specifications over unbounded domains. *In Proc. of FMCAD'10*. IEEE, 2010.
10. Topnik, Ch., Wilhelm, E., Margaria, T., Steffen, B.: jMosel: A stand-alone tool and jABC plugin for M2L(Str). *In Proc. of SPIN'06*. LNCS 3925, Springer, 2006.
11. Ganzow, T., Kaiser, L.: New algorithm for weak monadic second-order logic on inductive structures. *In Proc. of CSL'10*. LNCS 6247, Springer, 2010.
12. Wies, T., Muñiz, M., Kuncak, V.: An efficient decision procedure for imperative tree data structures. *In Proc. of CADE'11*. LNCS 6803, Springer, 2011.
13. Doyen, L., Raskin, J.F.: Antichain algorithms for finite automata. *In Proc. of TACAS'10*. LNCS 6015, Springer, 2010.
14. Wulf, M.D., Doyen, L., Henzinger, T.A., Raskin, J.F.: Antichains: A new algorithm for checking universality of finite automata. *In Proc. of CAV'06*. LNCS 4144, Springer, 2006.
15. Abdulla, P.A., Chen, Y.F., Holík, L., Mayr, R., Vojnar, T.: When simulation meets antichains (on checking language inclusion of NFAs). *In Proc. of TACAS'10*. LNCS 6015 , Springer, 2010.
16. Bustan, D., Grumberg, O.: Simulation based minimization. *In Proc. of CADE'00*. LNCS 1831, 2000.
17. Bonchi, F., Pous, D.: Checking NFA equivalence with bisimulations up to congruence. *In Proc. of POPL'13*. ACM, 2013.
18. Comon, H., Dauchet, M., Gilleron, R., Löding, C., Jacquemard, F., Lugiez, D., Tison, S., Tommasi, M.: Tree Automata Techniques and Applications, 2008.
19. Büchi, J.R.: Weak second-order arithmetic and finite automata. Technical report, The University of Michigan (1959) Available at URL: http://hdl.handle.net/2027.42/3930, 2010.
20. Fiedor, T., Holík, L., Lengál, O., Vojnar, T.: dWiNA, 2014. Available from http://www.fit.vutbr.cz/research/groups/verifit/tools/dWiNA/.
21. Lengál, O., Šimáček, J., Vojnar, T.: VATA: A library for efficient manipulation of non-deterministic tree automata. *In Proc. of TACAS'12*. LNCS 7214, Springer, 2012.
22. Abdulla, P.A., Bouajjani, A., Holík, L., Kaati, L., Vojnar, T.: Computing simulations over tree automata: Efficient techniques for reducing tree automata. *In Proc. of TACAS'08*. LNCS 4963, Springer, 2008.
23. Habermehl, P., Holík, L., Rogalewicz, A., Šimáček, J., Vojnar, T.: Forest automata for verification of heap manipulation. Formal Methods in System Design **41**(1), 2012.
24. Bouajjani, A., Habermehl, P., Holík, L., Touili, T., Vojnar, T.: Antichain-based universality and inclusion testing over nondeterministic finite tree automata. *In Proc. of CIAA'08*. LNCS 5148, Springer, 2008.
25. Fiedor, T., Holík, L., Lengál, O., Vojnar, T.: Nested Antichains for WS1S. Technical report FIT-TR-2014-06. http://www.fit.vutbr.cz/~ilengal/pub/FIT-TR-2014-06.pdf.