

Fixpoint Guided Abstraction Refinement for Alternating Automata*

Pierre Ganty¹, Nicolas Maquet^{2**} and Jean-François Raskin²

¹ University of California, Los Angeles, USA

² Université Libre de Bruxelles (ULB), Belgium

Abstract. In this paper, we develop and evaluate two new algorithms for checking emptiness of alternating automata. These algorithms build on previous works. First, they rely on antichains to efficiently manipulate the state-spaces underlying the analysis of alternating automata. Second, they are abstract algorithms with built-in refinement operators based on techniques that exploit information computed by abstract fixed points (and not counter-examples as it is usually the case). The efficiency of our new algorithms is illustrated by experimental results.

1 Introduction

Alternating automata are a generalization of both nondeterministic and universal automata. In an alternating automaton, the transition relation is defined using positive Boolean formulas: disjunctions allow for the expression of nondeterministic transitions and conjunctions allow for the expression of universal transitions. The emptiness problem for alternating automata being PSPACE-COMplete [3], several computationally-hard automata-theoretic and model-checking problems can be reduced in polynomial time to the emptiness problem for those automata. It is thus very desirable to design efficient algorithms for checking emptiness of those automata. In this paper, we propose new algorithms for efficiently checking the emptiness problem for alternating automata over finite words. Those new algorithms combine two recent lines of research.

First, we use efficient techniques based on *antichains*, initially introduced in [6], to symbolically manipulate the state-spaces underlying the analysis of alternating automata. Antichain-based techniques have been applied to several problems in automata theory [6, 8, 9, 1] and for solving games of imperfect information [13]. Those techniques have also been applied with success to the satisfiability and model-checking of LTL specifications [8]. Our team has implemented these algorithms in a tool called ALASKA [7], which is available for download³.

Second, to apply this antichain technique to even larger instances of alternating automata, we instantiate a generic abstract-refinement method that we have proposed

* Work supported by the projects: (i) Quasimodo: “Quantitative System Properties in Model-Driven-Design of Embedded”, <http://www.quasimodo.aau.dk/>, (ii) Gasics: “Games for Analysis and Synthesis of Interactive Computational Systems”, <http://www.ulb.ac.be/di/gasics/>, (iii) Moves: “Fundamental Issues in Modelling, Verification and Evolution of Software”, <http://moves.ulb.ac.be>, a PAI program funded by the Federal Belgian Government, (iv) CFV (Federated Center in Verification) funded by the FNRS <http://www.ulb.ac.be/di/ssd/cfv/>.

** This author is supported by an FNRS-FRIA grant.

³ See <http://www.antichains.be>

in [5] and further developed in [11, 12]. This abstract-refinement method does not use counter-examples to refine inconclusive abstractions contrary to most of the methods presented and implemented in the literature, see for example [4]. Instead, our algorithm uses the entire information computed by the abstract analysis and combines it with information obtained by one application of a concrete predicate transformer. The algorithm presented in [5] is a generic solution that does not lead directly to efficient implementations. In particular, as shown in [11], in order to obtain an efficient implementation of this algorithm, we need to define a family of abstract domains on which abstract analysis can be effectively computed, as well as practical operators to refine the elements of this family of abstract domains. In this paper, we use the set of *partitions* of the locations of an alternating automaton to define the family of abstract domains. Those abstract domains and their refinement operators can be used both in *forward* and *backward* algorithms for checking emptiness of alternating automata.

To show the practical interest of these new algorithms, we have implemented them into the ALASKA tool. We illustrate the efficiency of our new algorithms on examples of alternating automata constructed from LTL specifications interpreted over finite words. With the help of those examples, we show that our algorithms are able to concentrate the analysis on important parts of the state-space and abstract away the less interesting parts *automatically*. This allows us to treat much larger instances than with the concrete forward or backward algorithms. We are confident that those new algorithms will allow us to solve problems of practical relevance that are currently out of reach of automatic methods.

2 Preliminaries

Alternating Automata. Let S be a set. We denote $\mathcal{B}^+(S)$ the set of *positive Boolean formulas* over S . Formally, $\mathcal{B}^+(S) ::= s \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2$, where $s \in S$. A valuation for a set of proposition S is encoded as a subset of S . For each formula $\phi \in \mathcal{B}^+(S)$ we write $\llbracket \phi \rrbracket \subseteq 2^S$ the set of valuations that satisfy ϕ ; as usual, $c \in \llbracket \phi \rrbracket$ is interpreted as the valuation that assigns “true” only to the variables in c . Let Σ be a finite alphabet. A finite word w is a finite sequence $w = \sigma_0 \sigma_1 \dots \sigma_{n-1}$ of letters from Σ . We write Σ^* the set of finite words over Σ . We now recall the definition of *alternating automata over finite words* (AFA for short).

Definition 1. An alternating finite automaton is a tuple $\langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$ where : $\text{Loc} = \{l_1, \dots, l_n\}$ is the set of locations; $\Sigma = \{\sigma_1, \dots, \sigma_m\}$ is the set of alphabet symbols; $q_0 \in \text{Loc}$ is the initial location; $\delta: \text{Loc} \times \Sigma \rightarrow \mathcal{B}^+(\text{Loc})$ is the transition function; and $F \subseteq \text{Loc}$ is the set of accepting locations.

As we will often manipulate sets of sets of locations in the sequel, we will refer to the inner sets as *cells*. Let $\text{Cells}(S) = 2^S$. A *cell* of an AFA with locations Loc is an element of $\text{Cells}(\text{Loc})$. A set of cells X is \subseteq -upward-closed (resp. \subseteq -downward-closed) if for all $c \in X$ and for all $c' \in \text{Cells}(\text{Loc})$ s.t. $c \subseteq c'$ (resp. $c' \subseteq c$), we have $c' \in X$. Instead of defining the traditional notion of runs for AFA, we define their semantics as a *directed graph*, the nodes of which are cells. Each edge in the cell graph is labeled by an alphabet symbol.

Definition 2. Let $A = \langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$, $\llbracket A \rrbracket = \langle V, E \rangle$ where: $V = \text{Cells}(\text{Loc})$ and $\langle c, \sigma, c' \rangle \in E$ iff $c' \in \llbracket \bigwedge_{l \in c} \delta(l, \sigma) \rrbracket$. A word $w = \sigma_1, \dots, \sigma_p$ is accepted by the automaton A iff there exists a path c_0, c_1, \dots, c_p of cells of V such that $q_0 \in c_0$, $c_p \in \text{Cells}(F)$ (the set of accepting cells), and $\forall i \in [1, \dots, p] : \langle c_{i-1}, \sigma_i, c_i \rangle \in E$.

In the sequel, we will consider $\llbracket A \rrbracket$ simply as the set of edges E of the cell graph and leave the set of vertices V implicit.

Predicate Transformers. We have defined the semantics of alternating automata as a directed graph of cells. To explore this graph, we use *predicate transformers*.

Definition 3. We consider the following predicate transformers (A is an AFA):

$$\begin{aligned} \text{post}_\sigma[A](X) &= \{c_2 \mid \exists \langle c_1, \sigma, c_2 \rangle \in \llbracket A \rrbracket : c_1 \in X\} & \text{post}[A](X) &= \bigcup_{\sigma \in \Sigma} \text{post}_\sigma[A](X) \\ \widetilde{\text{post}}_\sigma[A](X) &= \{c_2 \mid \forall \langle c_1, \sigma, c_2 \rangle \in \llbracket A \rrbracket : c_1 \in X\} & \widetilde{\text{post}}[A](X) &= \bigcap_{\sigma \in \Sigma} \widetilde{\text{post}}_\sigma[A](X) \\ \text{pre}_\sigma[A](X) &= \{c_1 \mid \exists \langle c_1, \sigma, c_2 \rangle \in \llbracket A \rrbracket : c_2 \in X\} & \text{pre}[A](X) &= \bigcup_{\sigma \in \Sigma} \text{pre}_\sigma[A](X) \\ \widetilde{\text{pre}}_\sigma[A](X) &= \{c_1 \mid \forall \langle c_1, \sigma, c_2 \rangle \in \llbracket A \rrbracket : c_2 \in X\} & \widetilde{\text{pre}}[A](X) &= \bigcap_{\sigma \in \Sigma} \widetilde{\text{pre}}_\sigma[A](X) \end{aligned}$$

Theorem 1. Let $A = \langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$ an AFA, $\bar{X} \equiv \text{Cells}(\text{Loc}) \setminus X$, $\mathcal{F} = \text{Cells}(F)$. $L(A) = \emptyset$ iff $(\mu x \cdot \text{post}[A](x) \cup \llbracket q_0 \rrbracket) \subseteq \bar{\mathcal{F}}$ iff $(\mu x \cdot \text{pre}[A](x) \cup \mathcal{F}) \subseteq \llbracket q_0 \rrbracket$.

The lattice of partitions. The heart of our abstraction scheme is to *partition* the set of locations Loc of an AFA, in order to build a smaller (hopefully more manageable) automaton. We recall the notion of partitions and some of their properties. Let \mathcal{P} be a partition of the set $S = \{l_1, \dots, l_n\}$ into k classes (called *blocks* in the sequel) $\mathcal{P} = \{b_1, \dots, b_k\}$. Partitions are classically ordered as follows: $\mathcal{P}_1 \preceq \mathcal{P}_2$ iff $\forall b_1 \in \mathcal{P}_1, \exists b_2 \in \mathcal{P}_2 : b_1 \subseteq b_2$. It is well known, see [2], that the set of partitions together with \preceq form a complete lattice where $\{\{l_1\}, \dots, \{l_n\}\}$ is the \preceq -minimal element, $\{\{l_1, \dots, l_n\}\}$ is the \preceq -maximal element and the greatest lower bound of two partitions \mathcal{P}_1 and \mathcal{P}_2 , noted $\mathcal{P}_1 \wedge \mathcal{P}_2$, is the partition given by $\{b \neq \emptyset \mid \exists b_1 \in \mathcal{P}_1, \exists b_2 \in \mathcal{P}_2 : b = b_1 \cap b_2\}$. The least upper bound of two partitions \mathcal{P}_1 and \mathcal{P}_2 , noted $\mathcal{P}_1 \vee \mathcal{P}_2$, is the finest partition such that given $b \in \mathcal{P}_1 \vee \mathcal{P}_2$, for all $l_i \neq l_j : l_i \in b$ and $l_j \in b$ we have: $\exists b' \in \mathcal{P}_1 \vee \mathcal{P}_2 : l_i \in b'$ and $l_j \in b'$. Also, we shall use \mathcal{P} as a function such that $\mathcal{P}(l)$ simply returns the block b to which l belongs in \mathcal{P} .

3 Deciding AFA Emptiness Using Antichains

A fundamental problem regarding AFA is the *emptiness problem*; i.e., to decide if there exists at least one word accepted by an AFA. Since nondeterministic automata (NFA, for short) emptiness can be solved in linear-time, a natural solution is to first perform an AFA \rightarrow NFA translation and then check for emptiness. The translation is simple (albeit computationally difficult), as it amounts to a subset construction, similar to that of NFA determinization. Notice that the cell-graph semantics of AFA defined in the previous section is essentially an NFA obtained by subset construction. In earlier works [6, 8, 9], we have designed new efficient algorithms for AFA emptiness. Those algorithms are based on efficient manipulations of \subseteq -upward- or downward-closed sets of cells using *antichains*. In our context, an antichain is the unique set of \subseteq -minimal (resp. \subseteq -maximal) cells of an upward-closed (resp. downward-closed) set of cells X , which we denote by $\lfloor X \rfloor$ (resp. $\lceil X \rceil$). The crucial properties of antichains are that (i) they are canonical representations of \subseteq -closed sets of cells, (ii) the predicate transformers on AFA evaluate to \subseteq -closed sets (they can thus be canonically represented with antichains) and, (iii) evaluating a predicate transformer on any set of cells is equivalent to evaluating it on the \subseteq -closure of that set (we can thus evaluate predicate transformers *directly* on antichains, without losing any information). Due to lack of space, we do not recall the framework of antichains in this work (it can be found in [6]). In the sequel, we will assume that all the computations on sets of cells are performed using antichains.

4 Abstraction of Alternating Automata

4.1 Abstract domain

In this section, we present an original algorithmic framework for the analysis of AFA, using antichains along with abstract interpretation. Given an AFA with locations Loc , our algorithm will use a family of abstract domains defined by the set of partitions \mathcal{P} of Loc . The concrete domain is the complete lattice $2^{\text{Cells}(\text{Loc})}$, and each partition \mathcal{P} defines the abstract domain as $2^{\text{Cells}(\mathcal{P})}$. We refer to elements of $\text{Cells}(\text{Loc})$ as *concrete cells* and elements of $\text{Cells}(\mathcal{P})$ as *abstract cells*. An abstract cell is thus a set of blocks of the partition \mathcal{P} and it represents all the concrete cells which can be constructed by choosing at least one location from each block. To capture this representation role of abstract cells, we define the following predicate.

Definition 4. *The predicate $\text{Covers} : \text{Cells}(\mathcal{P}) \times \text{Cells}(\text{Loc}) \rightarrow \{\top, \perp\}$ is defined as follows: $\text{Covers}(c^\alpha, c)$ iff $c^\alpha = \{\mathcal{P}(l) \mid l \in c\}$.*

Example 1. Let $\text{Loc} = \{1, \dots, 5\}$, $\mathcal{P} = \{a : \{1\}, b : \{2, 3\}, c : \{4, 5\}\}$. We have that $\text{Covers}(\{a, c\}, \{1, 3\}) = \perp$, $\text{Covers}(\{a, c\}, \{1, 4\}) = \top$, and $\text{Covers}(\{a, c\}, \{1\}) = \perp$.

To make proper use of the theory of abstract interpretation, we define an *abstraction* and a *concretization* functions, and show that they form a *Galois connection* between the concrete domain and each of our abstract domains.

Definition 5. *Let \mathcal{P} be a partition of the set Loc , we define the functions*

$\alpha_{\mathcal{P}} : 2^{\text{Cells}(\text{Loc})} \rightarrow 2^{\text{Cells}(\mathcal{P})}$ and $\gamma_{\mathcal{P}} : 2^{\text{Cells}(\mathcal{P})} \rightarrow 2^{\text{Cells}(\text{Loc})}$ as follows :
 $\alpha_{\mathcal{P}}(X) = \{c^\alpha \mid \exists c \in X : \text{Covers}(c^\alpha, c)\}$, $\gamma_{\mathcal{P}}(X) = \{c \mid \exists c^\alpha \in X : \text{Covers}(c^\alpha, c)\}$.

Lemma 1. *For any partition \mathcal{P} of $\text{Loc} : (2^{\text{Cells}(\text{Loc})}, \subseteq) \xleftrightarrow[\alpha]{\gamma} (2^{\text{Cells}(\mathcal{P})}, \subseteq)$.*

In the sequel, we will omit the \mathcal{P} subscript of α and γ when the partition is clear from the context. Additionally, we define $\mu_{\mathcal{P}} = \gamma_{\mathcal{P}} \circ \alpha_{\mathcal{P}}$.

4.2 Efficient abstract analysis

In the sequel, we will need to evaluate fixpoint-expressions over the abstract domain. In theory, we could simply surround every predicate transformer occurring in the fixpoint-expressions by $\alpha \circ \cdot \circ \gamma$ to obtain an abstract fixpoint. However, for obvious performance concerns, we want to avoid as many concretization and abstraction steps as possible, and ideally make all the computations *directly over the abstract domain*. Furthermore, we would like that these *abstract predicate transformers* enjoy the same useful properties w.r.t. antichains so that we can reuse the results of the previous section. To achieve this goal, we proceed as follows. Given a partition \mathcal{P} of the set of locations of an alternating automaton, we use a *syntactic transformation* θ that builds an *abstract AFA* which over-approximates the behavior of the original automaton. Later in this section we will show that the *pre* and *post* predicate transformers can be directly evaluated on this abstract automaton to obtain the same result (but much faster) than the $\alpha \circ \cdot \circ \gamma$ computation on the original automaton. To express this syntactic transformation, we define *syntactic variants* of the abstraction and concretization functions.

Definition 6. Let \mathcal{P} be a partition of the set Loc . We define the following syntactic abstraction and concretization functions over positive Boolean formulas: $\hat{\alpha} : \mathcal{B}^+(\text{Loc}) \rightarrow \mathcal{B}^+(\mathcal{P})$ and $\hat{\gamma} : \mathcal{B}^+(\mathcal{P}) \rightarrow \mathcal{B}^+(\text{Loc})$, such that $\hat{\alpha}(l) = \mathcal{P}(l)$, $\hat{\alpha}(\phi_1 \vee \phi_2) = \hat{\alpha}(\phi_1) \vee \hat{\alpha}(\phi_2)$, and $\hat{\alpha}(\phi_1 \wedge \phi_2) = \hat{\alpha}(\phi_1) \wedge \hat{\alpha}(\phi_2)$. Likewise, $\hat{\gamma}(b) = \bigvee_{l \in b} l$, $\hat{\gamma}(\phi_1 \vee \phi_2) = \hat{\gamma}(\phi_1) \vee \hat{\gamma}(\phi_2)$, and $\hat{\gamma}(\phi_1 \wedge \phi_2) = \hat{\gamma}(\phi_1) \wedge \hat{\gamma}(\phi_2)$.

Lemma 2. For every $\phi \in \mathcal{B}^+(\text{Loc})$ we have that $\llbracket \hat{\alpha}(\phi) \rrbracket = \alpha(\llbracket \phi \rrbracket)$, and for every $\phi \in \mathcal{B}^+(\mathcal{P})$ we have that $\llbracket \hat{\gamma}(\phi) \rrbracket = \gamma(\llbracket \phi \rrbracket)$.

Definition 7. Let $A = \langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$ and \mathcal{P} a partition of Loc . $\theta(A, \mathcal{P}) = \langle \text{Loc}^\alpha, \Sigma, b_0, \delta^\alpha, F^\alpha \rangle$ where: $\text{Loc}^\alpha = \mathcal{P}$, $b_0 = \mathcal{P}(q_0)$, $\delta^\alpha(b, \sigma) = \hat{\alpha}(\bigvee_{l \in b} \delta(l, \sigma))$, and $F^\alpha = \{b \in \mathcal{P} \mid b \cap F \neq \emptyset\}$.

Theorem 2. Let A be an AFA, \mathcal{P} a partition of its locations and $A^\alpha = \theta(A, \mathcal{P})$, $\alpha \circ \text{post}[A] \circ \gamma = \text{post}[A^\alpha]$ and $\alpha \circ \text{pre}[A] \circ \gamma = \text{pre}[A^\alpha]$.

This theorem is crucial for the practical efficiency of our algorithms. In our framework, the evaluation of an abstract fixpoint on a large automaton amounts to compute a concrete fixpoint on a smaller automaton that is easy to obtain (the θ transformation can be done in linear time).

4.3 Precision of the abstract domain

We now present some results about precision and representability in our family of abstract domains. In particular, for the automatic refinement of abstract domains, we will need an effective way of computing the *coarsest partition* which can represent an upward- or downward closed set of cells without loss of precision.

Definition 8. A set of cells $X \subseteq \text{Cells}(\text{Loc})$ is representable in the abstract domain $2^{\text{Cells}(\mathcal{P})}$ iff $\mu_{\mathcal{P}}(X) = X$ (recall that $\mu_{\mathcal{P}} = \gamma_{\mathcal{P}} \circ \alpha_{\mathcal{P}}$).

Lemma 3. Let $X \subseteq \text{Cells}(\text{Loc})$, let \mathcal{P}_1 and \mathcal{P}_2 be two partitions of Loc . If X is representable with \mathcal{P}_1 and representable with \mathcal{P}_2 , then X is representable with $\mathcal{P}_1 \vee \mathcal{P}_2$.

As the lattice of partition is a complete lattice, we have the following corollary.

Corollary 1. For all $X \subseteq \text{Cells}(\text{Loc})$, there exists a coarsest partition $\mathcal{P} = \bigvee \{\mathcal{P}' \mid \mu_{\mathcal{P}'}(X) = X\}$ such that $\mu_{\mathcal{P}}(X) = X$.

For upward- and downward-closed sets, we have an efficient way to compute this coarsest partition. We start with upward-closed sets. To obtain an algorithm, we use the notion of *neighbour list*. The neighbour list of a location l with respect to an upward-closed set X , which we write $\mathcal{N}_X(l)$ is the set of subsets of Loc along which l appears in $\lfloor X \rfloor$.

Definition 9. Let $X \subseteq \text{Cells}(\text{Loc})$ be an upward-closed set. The neighbour list of a location $l \in \text{Loc}$ w.r.t. X is the set $\mathcal{N}_X(l) = \{c \setminus \{l\} \mid c \in \lfloor X \rfloor, l \in c\}$.

The following lemma states that if two locations share the same neighbour lists w.r.t. an upward-closed set X , then they can be put in the same partition block and preserve the representability of X . Conversely, X cannot be exactly represented by any partition which puts into the same block two locations that have different neighbour lists.

Lemma 4. For any partition \mathcal{P} of Loc , for any upward-closed set X , the set X is representable in $2^{\text{Cells}(\mathcal{P})}$ iff $\forall l, l' \in \text{Loc} \cdot \mathcal{P}(l) = \mathcal{P}(l') \rightarrow \mathcal{N}_X(l) = \mathcal{N}_X(l')$.

Corollary 2. For all upward-closed sets $X \subseteq \text{Cells}(\text{Loc})$, the partition \mathcal{P} induced by the equivalence relation $l \sim l'$ iff $\mathcal{N}_X(l) = \mathcal{N}_X(l')$ is the coarsest partition that is able to represent X . Assuming that $\lfloor X \rfloor$ has been computed, this partition is computable in $O(n \log n)$ set comparisons, where n is the size of $\lfloor X \rfloor$.

5 Abstraction Refinement Algorithm

| | |
|--|---|
| <p>Input: $A = \langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$ Output: True iff $L(A) = \emptyset$</p> <pre> 1 $\mathcal{P}_0 \leftarrow \{F, \text{Loc} \setminus F\}$ 2 $Z_0 \leftarrow \overline{\text{Cells}(F)}$ 3 for i in $0, 1, 2, \dots$ do 4 $A_i^\alpha \leftarrow \theta(A, \mathcal{P}_i)$ 5 $A_i^\alpha = \langle \text{Loc}^\alpha, \Sigma, b_0, \delta^\alpha, F^\alpha \rangle$ 6 $I_i \leftarrow \llbracket b_0 \rrbracket$ 7 $R_i \leftarrow \mu x \cdot (I_i \cup \text{post}[A_i^\alpha](x)) \cap \alpha_{\mathcal{P}_i}(Z_i)$ 8 if $\text{post}[A_i^\alpha](R_i) \subseteq \alpha_{\mathcal{P}_i}(Z_i)$ then 9 return True 10 if $\llbracket q_0 \rrbracket \not\subseteq Z_i$ then 11 return False 12 $Z_{i+1} \leftarrow \gamma_{\mathcal{P}_i}(R_i) \cap \widetilde{\text{pre}}[A](\gamma_{\mathcal{P}_i}(R_i))$ 13 $\mathcal{P}_{i+1} \leftarrow \bigvee \{\mathcal{P} \mid \mu_{\mathcal{P}}(Z_{i+1}) = Z_{i+1}\}$ </pre> | <p>Input: $A = \langle \text{Loc}, \Sigma, q_0, \delta, F \rangle$ Output: True iff $L(A) = \emptyset$</p> <pre> 1 $\mathcal{P}_0 \leftarrow \{\{q_0\}, \text{Loc} \setminus \{q_0\}\}$ 2 $Z_0 \leftarrow \llbracket q_0 \rrbracket$ 3 for i in $0, 1, 2, \dots$ do 4 $A_i^\alpha \leftarrow \theta(A, \mathcal{P}_i)$ 5 $A_i^\alpha = \langle \text{Loc}^\alpha, \Sigma, b_0, \delta^\alpha, F^\alpha \rangle$ 6 $B_i \leftarrow \text{Cells}(F^\alpha)$ 7 $R_i \leftarrow \mu x \cdot (B_i \cup \text{pre}[A_i^\alpha](x)) \cap \alpha_{\mathcal{P}_i}(Z_i)$ 8 if $\text{pre}[A_i^\alpha](R_i) \subseteq \alpha_{\mathcal{P}_i}(Z_i)$ then 9 return True 10 if $\text{Cells}(F) \not\subseteq Z_i$ then 11 return False 12 $Z_{i+1} \leftarrow \gamma_{\mathcal{P}_i}(R_i) \cap \widetilde{\text{post}}[A](\gamma_{\mathcal{P}_i}(R_i))$ 13 $\mathcal{P}_{i+1} \leftarrow \bigvee \{\mathcal{P} \mid \mu_{\mathcal{P}}(Z_{i+1}) = Z_{i+1}\}$ </pre> |
|--|---|

Fig. 1. The *abstract-forward* (left) and *abstract-backward* (right) FGAR algorithms.

This section presents two fixpoint-guided abstraction refinement algorithms for AFA. These algorithms share several ideas with the generic algorithm presented in [5] but they are formally different, so we provide arguments showing their correctness. We concentrate here on explanations related to the abstract forward algorithm. The abstract backward algorithm is the dual of this algorithm and its correctness can be established in a very similar way. We first give an informal presentation of the ideas underlying the algorithm and then we expose formal arguments for its soundness and completeness.

Description of the forward abstract algorithm. The most important information computed in the algorithm is Z_i , which is an over-approximation of the set of reachable cells which cannot reach an accepting cell in i steps or less. In other words, all the cells outside Z_i are either unreachable, or can lead to an accepting cell in i steps or less (or both). Our algorithm always uses the coarsest partition \mathcal{P}_i that allows Z_i to be represented in the corresponding abstract domain. The algorithm begins by initializing Z_0 with the set of non-accepting cells and by initializing \mathcal{P}_0 accordingly (lines 1 and 2). The main loop proceeds as follows. First, we compute the abstract reachable cells R_i which are within Z_i , which is done by applying the θ transformation using \mathcal{P}_i (line 4), and by computing a forward abstract fixpoint (line 7). If R_i does not contain a cell which can leave Z_i , we know (as we will formally prove later in this section) that the

automaton is empty (line 8). If on the other hand, an initial cell (i.e., a cell containing q_0) is no longer in Z_i then we know that it can lead to an accepting cell in i steps or less (as it is obviously reachable) and we conclude that the automaton is non-empty (line 11). In the case where both tests failed, we *refine* the information contained in Z_i by removing all the cells which can leave R_i in one step, as we know that these cells are either surely unreachable or can lead to an accepting cell in $i + 1$ steps or less. Finally, the current abstract domain is changed to be able to represent the new Z_i (line 13), using the neighbour list algorithm of Corollary 2. It is important to note that this refinement operation is not the traditional refinement used in counter-example guided abstraction refinement. Note also that our algorithm does not necessarily choose a new abstract domain that is strictly more precise than the previous one as in [5]. Instead, the algorithm uses the most abstract domain possible at all times. As we cannot rely on the termination proof from [5], we provide a new one at the end of this section.

Completeness and correctness of the forward abstract algorithm. Correctness and completeness relies on the properties formalized in the following lemma.

Lemma 5. *Let $\text{Reach} = \mu x \cdot \llbracket q_0 \rrbracket \cup \text{post}[A](x)$ be the reachable cells of A , let $\text{Bad}^k = \bigcup_{j=0}^{j=k} \text{pre}^j[A](\text{Cells}(F))$ be the cells that can reach an accepting cell in k steps or less, and let us note $\text{Safe}^k = \text{Cells}(\text{Loc}) \setminus \text{Bad}^k$, i.e. the set of cells that cannot reach an accepting cell in k steps or less. The following four properties hold:*

1. $\forall i \geq 0: \mu_{\mathcal{P}_i}(Z_i) = Z_i$, i.e. Z_i is representable in the successive abstract domains;
2. $\forall i \geq 0: Z_{i+1} \subseteq Z_i$, i.e. the sets Z_i are decreasing;
3. $\forall i \geq 0: \text{Reach} \cap \text{Safe}^i \subseteq Z_i$, i.e. Z_i over-approximates the reachable cells that cannot reach an accepting cell in i steps or less;
4. if $Z_i = Z_{i+1}$ then $\text{post}[A_i^\alpha](R_i) \subseteq \alpha_{\mathcal{P}_i}(Z_i)$.

Theorem 3. *The forward abstract algorithm with refinement is sound and complete to decide the emptiness of AFA.*

6 Experimental Evaluation

In this section, we evaluate the practical performance of our techniques with three series of benchmarks. Each benchmark is composed of a pair of LTL formulas $\langle \psi, \phi \rangle$ interpreted on finite words, and for which we want to know if ϕ is a logical consequence of ψ , i.e. if $\psi \models \phi$ holds. To solve this problem, we translate the formula $\psi \wedge \neg\phi$ into an AFA and check that the language of the AFA is empty. This translation is linear in the size of the formula and creates a location in the AFA for each subformula. As we will see, our ψ formulas are constructed as large conjunctions of constraints and model the behavior of finite-state systems, while the ϕ formulas model properties of those systems. We defined properties with varying degrees of *locality*. Intuitively, a property ϕ is local when only a small number of subformulas of ψ are needed to establish $\psi \models \phi$. This is not a formal notion but it will be clear from the examples. We will show in this section that our abstract algorithms are able to automatically identify subformulas which are not needed to establish the property. Due to lack of space, we only report results where $\psi \models \phi$ holds. We now present each benchmark in turn.

Benchmark 1. The first benchmark takes 2 parameters $n > 0$ and $0 < k \leq n$: $\text{Bench1}(n, k) = \langle \bigwedge_{i=0}^{n-1} G(p_i \rightarrow (F(\neg p_i) \wedge F(p_{i+1}))), Fp_0 \rightarrow Fp_k \rangle$. Clearly we have that $\psi \models \phi$ holds for all values of k and also that the subformulas of ψ for $i > k$ are not needed to establish $\psi \models \phi$.

Benchmark 2. This second benchmark is used to demonstrate how our algorithms can automatically detect less obvious versions of locality than for Bench1. It uses 2 parameters k and n with $0 < k \leq n$ and is built using the following recursive nesting definition: $\text{Sub}(n, 1) = Fp_n$; for odd values of $k > 1$ $\text{Sub}(n, k) = F(p_n \wedge X(\text{Sub}(n, k - 1)))$; and for even values of $k > 1$ $\text{Sub}(n, k) = F(\neg p_n \wedge X(\text{Sub}(n, k - 1)))$. Our second benchmark is : $\text{Bench2}(n, k) = \langle \bigwedge_{i=0}^{n-1} G(p_i \rightarrow \text{Sub}(i + 1, k)), Fp_0 \rightarrow Fp_n \rangle$. It is relatively easy to see that $\psi \models \phi$ holds for any value of k , and that for odd values of k , the nested subformulas beyond the first level are not needed to establish the property.

Benchmark 3. This third and final benchmark aims to demonstrate the usefulness of our abstraction algorithms in a more realistic setting. We specified the behavior of a lift with n floors with a parametric LTL formula. For n floors, $\text{Prop} = \{f_1, \dots, f_n, b_1, \dots, b_n, \text{open}\}$. The f_i propositions represent the current floor. Only one of the f_i 's can be true at any time, which is initially f_1 . The b_i propositions represent the state (lit or unlit) of the call-buttons of each floor and there is only one button per floor. The additional *open* proposition is true when the doors of the lift are open. The constraints on the dynamics of this system are as follows : (i) initially the lift is at the first floor and the doors are open, (ii) the lift must close its doors when changing floors, (iii) the lift must go through floors in the correct order, (iv) when a button is lit, the lift eventually reaches the corresponding floor and opens its doors, and finally (v) when the lift reaches a floor, the corresponding button becomes unlit. Let n be the number of floors. We apply our algorithms to check two properties which depend on a parameter k with $1 < k \leq n$, namely $\text{Spec1}(k) = G((f_1 \wedge b_k) \rightarrow (\neg f_k U f_{k-1}))$, and $\text{Spec2}(k) = G((f_1 \wedge b_k \wedge b_{k-1}) \rightarrow (b_k U \neg b_{k-1}))$.

Experimental results. All the results of our experiments are found in Fig. 2, and were performed on a quad-core 3,2 Ghz Intel CPU with 12 Gb of memory. Due to lack of space, we only report results for the concrete forward and abstract backward algorithms which were the fastest (by a large factor) in all our experiments. The columns of the table are as follows. *ATC* is the size of the largest antichain encountered, *iters* is the number of iterations of the fixpoint in the concrete case and the maximal number of iterations of all the abstract fixpoints in the abstract case, ATC^α and ATC^γ are respectively the sizes of the largest abstract and concrete antichains encountered, *steps* is the number of execution of the refinement steps and $|\mathcal{P}|$ is the maximum number of blocks in the partitions.

Benchmark 1. The partition sizes of the first benchmark illustrate how our algorithm exploits the locality of the property to abstract away the irrelevant parts of the system. For local properties, i.e. for small values of k , $|\mathcal{P}|$ is small compared to $|\text{Loc}|$ meaning that the algorithm automatically ignores many subformulas which are irrelevant to the property. For larger values of k , the abstraction overhead becomes larger, but that overhead becomes less important as the system grows.

Benchmark 2. On the second benchmark, our abstract algorithm largely outperforms the concrete algorithm. Notice how for $k \geq 3$ the partition sizes do not continue to grow (it also holds for values of k beyond 5). This means that contrary to the concrete algorithm, FGAR does not get trapped in the intricate nesting of the F modalities (which are not

necessary to prove the property) and abstracts it completely with a constant number of partition blocks. The speed improvement is considerable.

Benchmark 3. On this final benchmark, the abstract algorithm outperforms the concrete algorithm when the locality of the property spans less than 5 floors. Beyond that value, the abstract algorithm starts to take longer than the concrete version. From the *ATC* column, the antichain sizes remain constant in the concrete algorithm, when the number of floors increases. This indicates that the difficulty of this benchmark comes mainly from the exponential size of the alphabet rather than the state-space itself. As our algorithms only abstracts the locations and not the alphabet, these results are not surprising.

7 Discussion

We have proposed in this paper two new abstract algorithms with refinement for deciding language emptiness for AFA. Our algorithm is based on an abstraction-refinement scheme inspired from [5], which is different from the usual refinement techniques based on counter-example elimination [4]. Our algorithm also builds on the successful technique of antichains, that we have introduced in [6], to symbolically manipulate closed sets of cells (sets of sets of locations). We have demonstrated with a set of benchmarks that our algorithm is able to find coarse abstractions for complex automata constructed from large LTL formulas. For a large number of instances of those benchmarks, the abstract algorithms outperform by several order of magnitude the concrete algorithms. We believe that this clearly shows the interest of our new algorithms and their potential future developments.

Acknowledgments The authors would like to thank Gilles Geeraerts for some fruitful discussions on the abstraction scheme.

References

1. A. Bouajjani, P. Habermehl, L. Holík, T. Touili, and T. Vojnar. Antichain-based universality and inclusion testing over nondeterministic finite tree automata. In *CIAA*, pages 57–67, 2008.
2. S. Burris and H. P. Sankappanavar. *A Course in Universal Algebra*. Springer, 1981.
3. A. K. Chandra, D. Kozen, and L. J. Stockmeyer. Alternation. *J. ACM*, 28(1):114–133, 1981.
4. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.
5. P. Cousot, P. Ganty, and J.-F. Raskin. Fixpoint-guided abstraction refinements. In *SAS*, volume 4634 of *LNCS*, pages 333–348. Springer, 2007.
6. M. De Wulf, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Antichains: A new algorithm for checking universality of finite automata. In *CAV*, volume 4144 of *LNCS*, pages 17–30, 2006.
7. M. De Wulf, L. Doyen, N. Maquet, and J.-F. Raskin. Alaska. In *ATVA*, pages 240–245, 2008.
8. M. De Wulf, L. Doyen, N. Maquet, and J.-F. Raskin. Antichains: Alternative algorithms for LTL satisfiability and model-checking. In *TACAS*, volume 4963 of *LNCS*, 2008.
9. L. Doyen and J.-F. Raskin. Improved algorithms for the automata-based approach to model-checking. In *TACAS*, volume 4424 of *LNCS*, pages 451–465. Springer, 2007.
10. S. Fogarty and M. Vardi. Buechi complementation and size-change termination. In *TACAS*, volume 5505 of *LNCS*, pages 16–30, 2009.
11. P. Ganty. *The Fixpoint Checking Problem: An Abstraction Refinement Perspective*. PhD thesis, Université Libre de Bruxelles, 2007.
12. P. Ganty, J.-F. Raskin, and L. Van Begin. From many places to few: automatic abstraction refinement for petri nets. *Fundamenta Informaticae*, 88(3):275–305, 2008.
13. J.-F. Raskin, K. Chatterjee, L. Doyen, and T. A. Henzinger. Algorithms for omega-regular games with imperfect information. *Logical Methods in Computer Science*, 3(3), 2007.

| | | | | <i>concrete forward</i> | | | <i>abstract backward</i> | | | | | | | |
|--------------|--------------|-----|------|-------------------------|--------|--------|--------------------------|------------------|------------------|-------|-------|-----------------|----|----|
| n | k | Loc | Prop | time | ATC | iters | time | ATC ^α | ATC ^γ | iters | steps | $ \mathcal{P} $ | | |
| Bench1 | 11 | 5 | 50 | 12 | 0,10 | 6 | 3 | 0,23 | 55 | 2 | 5 | 3 | 27 | |
| | 15 | 5 | 66 | 16 | 1,60 | 6 | 3 | 0,56 | 55 | 2 | 5 | 3 | 31 | |
| | 19 | 5 | 82 | 20 | 76,62 | 6 | 3 | 8,64 | 55 | 2 | 5 | 3 | 35 | |
| | 11 | 7 | 50 | 12 | 0,13 | 8 | 3 | 0,87 | 201 | 2 | 5 | 3 | 31 | |
| | 15 | 7 | 66 | 16 | 2,04 | 8 | 3 | 1,21 | 201 | 2 | 5 | 3 | 35 | |
| | 19 | 7 | 82 | 20 | 95,79 | 8 | 3 | 9,99 | 201 | 2 | 5 | 3 | 39 | |
| | 11 | 9 | 50 | 12 | 0,16 | 10 | 3 | 12,60 | 779 | 2 | 5 | 3 | 35 | |
| | 15 | 9 | 66 | 16 | 2,69 | 10 | 3 | 13,42 | 779 | 2 | 5 | 3 | 39 | |
| | 19 | 9 | 82 | 20 | 125,85 | 10 | 3 | 46,47 | 779 | 2 | 5 | 3 | 43 | |
| | Bench2 | 7 | 1 | 19 | 8 | 0,06 | 8 | 2 | 0,10 | 11 | 2 | 4 | 3 | 14 |
| | | 10 | 1 | 25 | 11 | 0,06 | 10 | 2 | 0,10 | 14 | 2 | 4 | 3 | 17 |
| | | 13 | 1 | 31 | 14 | 0,08 | 14 | 2 | 0,12 | 17 | 2 | 4 | 3 | 20 |
| 7 | | 3 | 33 | 8 | 0,78 | 201 | 14 | 0,13 | 11 | 2 | 4 | 3 | 26 | |
| 10 | | 3 | 45 | 11 | 802,17 | 4339 | 20 | 0,30 | 14 | 2 | 4 | 3 | 35 | |
| 13 | | 3 | 57 | 14 | > 1000 | - | - | 1,26 | 17 | 2 | 4 | 3 | 44 | |
| 7 | | 5 | 47 | 8 | 88,15 | 2122 | 26 | 0,14 | 11 | 2 | 4 | 3 | 26 | |
| 10 | | 5 | 65 | 11 | > 1000 | - | - | 0,37 | 14 | 2 | 4 | 3 | 35 | |
| 13 | | 5 | 83 | 14 | > 1000 | - | - | 1,47 | 17 | 2 | 4 | 3 | 44 | |
| Lift : Spec1 | | 8 | 3 | 84 | 17 | 0,30 | 10 | 17 | 0,51 | 23 | 40 | 7 | 4 | 21 |
| | | 12 | 3 | 116 | 25 | 17,45 | 10 | 25 | 1,63 | 23 | 40 | 7 | 4 | 21 |
| | | 16 | 3 | 148 | 33 | 498,65 | 10 | 33 | 26,65 | 23 | 40 | 7 | 4 | 21 |
| | 8 | 4 | 84 | 17 | 0,26 | 10 | 17 | 1,29 | 37 | 72 | 10 | 6 | 24 | |
| | 12 | 4 | 116 | 25 | 17,81 | 10 | 25 | 5,02 | 37 | 72 | 10 | 6 | 24 | |
| | 16 | 4 | 148 | 33 | 555,44 | 10 | 33 | 78,75 | 37 | 72 | 10 | 6 | 24 | |
| | 8 | 5 | 84 | 17 | 0,32 | 10 | 17 | 3,70 | 42 | 141 | 12 | 8 | 27 | |
| | 12 | 5 | 116 | 25 | 20,24 | 10 | 25 | 47,45 | 42 | 141 | 12 | 8 | 27 | |
| | 16 | 5 | 148 | 33 | 543,27 | 10 | 33 | > 1000 | - | - | - | - | - | |
| | Lift : Spec2 | 8 | 3 | 84 | 17 | 0,46 | 10 | 17 | 1,18 | 58 | 72 | 8 | 4 | 22 |
| | | 12 | 3 | 116 | 25 | 17,98 | 10 | 25 | 3,64 | 58 | 72 | 8 | 4 | 22 |
| | | 16 | 3 | 148 | 33 | 557,75 | 10 | 33 | 48,90 | 58 | 72 | 8 | 4 | 22 |
| 8 | | 4 | 84 | 17 | 0,29 | 10 | 17 | 3,04 | 124 | 126 | 11 | 6 | 25 | |
| 12 | | 4 | 116 | 25 | 19,29 | 10 | 25 | 10,63 | 124 | 126 | 11 | 6 | 25 | |
| 16 | | 4 | 148 | 33 | 576,56 | 10 | 33 | 128,40 | 124 | 126 | 11 | 6 | 25 | |
| 8 | | 5 | 84 | 17 | 0,31 | 10 | 17 | 15,88 | 131 | 266 | 14 | 8 | 28 | |
| 12 | | 5 | 116 | 25 | 19,47 | 10 | 25 | 283,90 | 131 | 266 | 14 | 8 | 28 | |
| 16 | | 5 | 148 | 33 | 568,83 | 10 | 33 | > 1000 | - | - | - | - | - | |

Fig. 2. Experimental results. Times are in seconds.