Mediating for Reduction
(on Minimizing Alternating Büchi Automata)

Parosh A. Abdulla[1], Yu-Fang Chen[1], Lukáš Holík[2], Tomáš Vojnar[2]

[1] University of Uppsala, Sweden,
{parosh,yu-fang.chen}@it.uu.se

[2] FIT, Brno University of Technology, Czech Republic,
{holik,vojnar}@fit.vutbr.cz

ABSTRACT. We propose a new approach for minimizing alternating Büchi automata (ABA). The approach is based on the so called *mediated equivalence* on states of ABA, which is the maximal equivalence contained in the so called *mediated preorder*. Two states p and q can be related by the mediated preorder if there is a *mediator* (mediating state) which forward simulates p and backward simulates q. Under some further conditions, letting a computation on some word jump from q to p (due to they get collapsed) preserves the language as the automaton can anyway already accept the word without jumps by runs through the mediator. We further show how the mediated equivalence can be computed efficiently. Finally, we show that, compared to the standard forward simulation equivalence, the mediated equivalence can yield much more significant reductions when applied within the process of complementing Büchi automata where ABA are used as an intermediate model.

1 Introduction

Alternating Büchi automata (ABA) are succinct state-machine representations of ω-regular languages (regular sets of infinite sequences). They are widely used in the area of formal specification and verification of non-terminating systems. One of the most prominent examples of the use of ABA is the complementation of nondeterministic Büchi automata [9]. It is an essential step of the automata-theoretic approach to model checking when the specification is given as a positive Büchi automaton [12] and also learning based model checking for liveness properties [4]. The other important usage of ABA is as the intermediate data structure for translating a linear temporal logic (LTL) specification to an automaton [7].

However, because of the compactness of ABA*, usually the algorithms that work on them are of high complexity. For example, both the complementation and the LTL translation algorithms transform an intermediate ABA to an equivalent NBA. The transformation is exponential in the size of the input ABA. Hence, one may prefer to reduce the size of the ABA (with some relatively cheaper algorithm) before giving it to the exponential procedure.

In the study of Fritz and Wilke, simulation-based minimization is proven as a very effective tool for reducing the size of ABA [6]. However, they considered only *forward* simulation relations. Inspired by some previous works [1], we believe that *backward* simulation can be used for reducing the size of ABA as well. Unfortunately, as will be explained in Section 3, quotienting wrt. *backward* simulation (i.e., simplify the automaton by collapsing backward simulation equivalent states) does not preserve the language.

*ABA's are exponentially more succinct than the nondeterministic ones.

In this paper, we develop an approach that uses backward simulation for simplifying ABA indirectly. Instead of looking for a suitable fragment of backward simulation that can be used to reduce the number of states of an ABA, we combine backward and forward simulation to form an even coarser relation called *mediated preorder* that can be used for minimization. The performance of minimizing ABA with *mediated preorder* is evaluated on a large set of experiments. In the experiments, we apply different simulation-based minimization approaches to improve the complementation algorithm of nondeterministic Büchi automata. The experimental results show that the minimization using mediated preorder significantly outperforms the minimization using forward simulation. To be more specific, in average, mediated minimization results in a 30% better reduction in the number of states and 50% better reduction in the number of transitions than forward minimization on the intermediate ABA. Moreover, in the complemented nondeterministic Büchi automata, mediated minimization results in a 100% better reduction in the number of states and 300% better reduction in the number of transitions than forward minimization.

2 Basic Definitions

Given a finite set X, we use X^* to denote the set of all finite words over X and X^ω for the set of all infinite words over X. The empty word is denoted ϵ and $X^+ = X^* \setminus \{\epsilon\}$. The concatenation of a finite word $u \in X^*$ and a finite or infinite word $v \in X^* \cup X^\omega$ is denoted by uv. For a word $w \in X^* \cup X^\omega$, $|w|$ is the length of w ($|w| = \infty$ if $w \in X^\omega$), w_i is the ith letter of w and w^i the ith prefix of w (the word u with $w = uv$ and $|u| = i$). $w^0 = \epsilon$. The concatenation of a finite word u and a set $S \subseteq X^* \cup X^\omega$ is defined as $uS = \{uv \mid v \in S\}$.

An *alternating Büchi automaton* is a tuple $\mathcal{A} = (\Sigma, Q, \iota, \delta, \alpha)$ where Σ is a finite alphabet, Q is a finite set of states, $\iota \in Q$ is an initial state, $\alpha \subseteq Q$ is a set of accepting states, and $\delta : Q \times \Sigma \to 2^{2^Q}$ is a total transition function. A *transition* of $\mathcal{A}$ is of the form $p \xrightarrow{a} P$ where $P \in \delta(q, a)$.

A *tree T over Q* is a subset of Q^+ that contains all nonempty prefixes of each one of its elements (i.e., $T \cup \{\epsilon\}$ is prefix-closed). Furthermore, we require that T contains exactly one $r \in Q$, the *root of T*, denoted *root*(T). We call the elements of Q^+ *paths*. For a path πq, we use *leaf*(πq) to denote its last element q. Define the set *branches*$(T) \subseteq Q^+ \cup Q^\omega$ such that $\pi \in$ *branches*(T) iff T contains all prefixes of π and π is not a proper prefix of any path in T. In other words, a *branch of T* is either a maximal path of T, or it is a word from Q^ω such that T contains all its nonempty prefixes. We use $succ_T(\pi) = \{r \mid \pi r \in T\}$ to denote the set of successors of a path π in T, and *height*(T) to denote the length of the longest branch of T. The tree U over Q is a *prefix of T* iff $U \subseteq T$ and for every $\pi \in U$, $succ_U(\pi) = succ_T(\pi)$ or $succ_U(\pi) = \emptyset$. The *suffix of T* defined by a path πq is the tree $T(\pi q) = \{q\psi \mid \pi q \psi \in T\}$.

Given a word $w \in \Sigma^\omega$, a tree T over Q is a *run of $\mathcal{A}$ on w*, if for every $\pi \in T$, *leaf*$(\pi) \xrightarrow{w_{|\pi|}} succ_T(\pi)$ is a transition of $\mathcal{A}$. Finite prefixes of T are called *partial runs on w*. A run T of $\mathcal{A}$ over w is *accepting* iff every infinite branch of T contains infinitely many accepting states. A word w is *accepted* by $\mathcal{A}$ from a state $q \in Q$ iff there exists an accepting run T of $\mathcal{A}$ over w with *root*$(T) = q$. The *language of a state $q \in Q$ in $\mathcal{A}$*, denoted $\mathcal{L}_\mathcal{A}(q)$, is the set of all words accepted by $\mathcal{A}$ from q. Then $\mathcal{L}(\mathcal{A}) = \mathcal{L}_\mathcal{A}(\iota)$ is the *language of $\mathcal{A}$*. For simplicity of presentation, we assume in the rest of the paper that δ never allows a transition of the form $p \xrightarrow{a} \emptyset$. This means that no run can contain a finite branch. Any automaton can be easily transformed into one without such transitions by adding a new accepting state q with $\delta(q, a) = \{\{q\}\}$ for every $a \in \Sigma$ and replacing every transition $p \xrightarrow{a} \emptyset$ by $p \xrightarrow{a} \{q\}$.

3 Simulation Relations

In this section, we give the definitions of forward and backward simulation over ABA and discuss some of their properties. The notion of backward simulation is inspired by a similar tree automata notion studied in [1, 3]—namely, the upward simulation parametrised by a downward simulation (the connection between tree automata and ABA follows from the fact that the runs of ABA are in fact trees).

For the rest of the section, we fix an ABA $\mathcal{A} = (\Sigma, Q, \iota, \delta, \alpha)$. We define relations $\preceq_\alpha$ and $\preceq_\iota$ on Q s.t. $q \preceq_\alpha r$ iff $q \in \alpha \implies r \in \alpha$ and $q \preceq_\iota r$ iff $q = \iota \implies r = \iota$. For a binary relation $\preceq$ on a set X, the relation $\preceq^{\forall\exists}$ on subsets of X is defined as $Y \preceq^{\forall\exists} Z$ iff $\forall z \in Z.\, \exists y \in Y.\, y \preceq z$, i.e., iff the upward closure of Z wrt. $\preceq$ is a subset of the upward closure of Y wrt. $\preceq$.

Forward Simulation. A *forward simulation* on $\mathcal{A}$ is a relation $\preceq_F \subseteq Q \times Q$ such that $p \preceq_F r$ implies that (i) $p \preceq_\alpha r$ and (ii) for all $p \xrightarrow{a} P$, there exists a $r \xrightarrow{a} R$ such that $P \preceq_F^{\forall\exists} R$.

For the basic properties of forward simulation, we rely on the work [8] by Gurumurthy et al. In particular, (i) there exists a unique maximal forward simulation $\preceq_F$ on $\mathcal{A}$ which is reflexive and transitive, (ii) for any $q, r \in Q$ such that $q \preceq_F r$, it holds that $\mathcal{L}_\mathcal{A}(q) \subseteq \mathcal{L}_\mathcal{A}(r)$, and (iii) quotienting wrt. $\preceq_F \cap \preceq_F^{-1}$ preserves the language of $\mathcal{A}$.

Backward Simulation. Let $\preceq_F$ be a forward simulation on $\mathcal{A}$. A *backward simulation* on $\mathcal{A}$ parameterized by $\preceq_F$ is a relation $\preceq_B \subseteq Q \times Q$ such that $p \preceq_B r$ implies that (i) $p \preceq_\iota r$, (ii) $p \preceq_\alpha r$, and (iii) for all $q \xrightarrow{a} P \cup \{p\}, p \notin P$, there exists a $s \xrightarrow{a} R \cup \{r\}, r \notin R$ such that $q \preceq_B s$ and $P \preceq_F^{\forall\exists} R$. The below lemma describes some properties of backward simulation.

LEMMA 1. *For any reflexive and transitive forward simulation $\preceq_F$ on $\mathcal{A}$, there exists a unique maximal backward simulation $\preceq_B$ on $\mathcal{A}$ parameterized by $\preceq_F$ that is reflexive and transitive.*

Backward simulation itself cannot be used for quotienting. In Appendix F.1, we give an example of an automaton, where quotienting using backward simulation does not preserve language. However, in Section 4.1, we show how backward simulation can be used to define a new relation for reducing ABA.

Let $\preceq_F$ and $\preceq_B$ be forward and backward simulations on $\mathcal{A}$, which are both reflexive and transitive. For every $x \in \{B, F, \alpha\}$, we extend the relation $\preceq_x$ to $Q^+ \times Q^+$ such that for $\pi, \psi \in Q^+$, $\pi \preceq_x \psi$ iff $|\pi| = |\psi|$ and for all $1 \leq i \leq |\pi|$, $\pi_i \preceq_x \psi_i$. We say that ψ forward simulates π, ψ backward simulates π, or ψ is more accepting than π when $\pi \preceq_F \psi$, $\pi \preceq_B \psi$, or $\pi \preceq_\alpha \psi$, respectively. This notation is further extended to trees. For trees T, U over Q and for $x \in \{\alpha, F\}$, we write, $T \preceq_x U$ if $branches(T) \preceq_x^{\forall\exists} branches(U)$. Similarly, we say that U forward simulates T, or U is more accepting than T when $T \preceq_F U$, or $T \preceq_\alpha U$, respectively. Note that $\preceq_x$ is reflexive and transitive for all the variants of $x \in \{F, B, \alpha\}$ defined over states, paths, or trees (this follows from the assumption that the original relations $\preceq_F$ and $\preceq_B$ on states are reflexive and transitive). Moreover, $\preceq_B \subseteq \preceq_\alpha$, $\preceq_B \subseteq \preceq_\iota$, and $\preceq_F \subseteq \preceq_\alpha$.

The following two lemmas formulate properties of the simulation relations that we will use in the rest of the paper.

LEMMA 2. *For any $p, r \in Q$ with $p \preceq_F r$ and a partial run T of $\mathcal{A}$ on $w \in \Sigma^\omega$ with the root p, there is a partial run U of $\mathcal{A}$ on w with the root r such that $T \preceq_F U$.*

For a tree T over Q, $\pi \in T$, and $1 \leq i \leq |\pi|$, the set $T \ominus_i \pi$ is the union of branches of suffix trees $T(\pi^i q), q \in succ_T(\pi^i)$, with the branches of the suffix tree $T(\pi^{i+1})$ excluded.

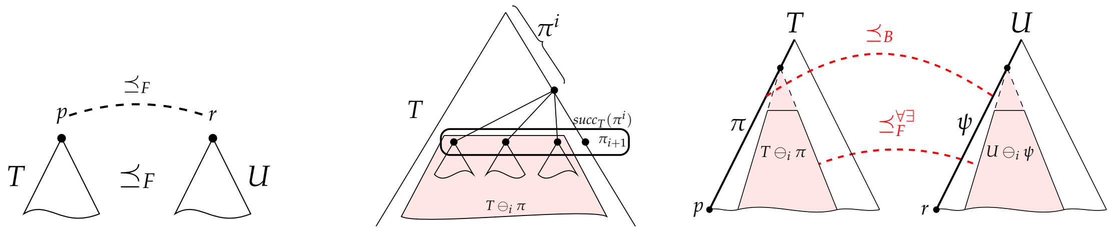

(a) Lemma 2 (b) The notation $T \ominus_i \pi$ (c) Lemma 3

Figure 1: Illustration of the lemmas

Formally, let $Q^i = succ_T(\pi^i) \setminus \{\pi_{i+1}\}$ be the set of all successors of π^i in T without the successor continuing in π. Then $T \ominus_i \pi = \bigcup_{q \in Q^i} branches(T(\pi^i q))$ (notice that if $i = 0$, then $T \ominus_i \pi = \emptyset$).

Lemma 3. *For any $p, r \in Q$ with $p \preceq_B r$, a partial run T of $\mathcal{A}$ on $w \in \Sigma^\omega$ and $\pi \in branches(T)$ with $leaf(\pi) = p$, there is a partial run U of $\mathcal{A}$ on w and $\psi \in branches(U)$ with $leaf(\psi) = r$ such that $\pi \preceq_B \psi$, and for all $1 \leq i \leq |\pi|$, $T \ominus_i \pi \preceq_F^{\forall\exists} U \ominus_i \psi$.*

4 Mediated Equivalence and Quotienting

Here we discuss the possibility of an indirect use of backward simulation for simplifying ABA via quotienting. We do not look for a suitable fragment of backward simulation only. Instead, we (1) combine backward and forward simulation to form an equivalence that subsumes both backward and forward simulation equivalence and (2) take a certain fragment of this equivalence, called *mediated equivalence*, that can be used for quotienting.

4.1 The Notion and Intuition of Mediated Equivalence

Collapsing states of an automaton wrt. some equivalence allows a run that arrives to some state to *jump* to another equivalent state and continue from there. Alternatively, this can be viewed as *extending* the source state of the jump by the outgoing transitions of the target state[†]. The equivalence must have the property that the language is not increased even when the jumps (or, alternatively, transition extensions) are allowed. This is what we aim at when introducing the *mediated equivalence* $\equiv_M$ based on a so called *mediated preorder* $\preceq_M$. The mediated preorder $\preceq_M$ will in particular be defined as a suitable transitive fragment of $\preceq_F \circ \preceq_B^{-1}$ in the following.

The intuition behind allowing a run to jump from a state r to a state q such that $q \preceq_F \circ \preceq_B^{-1} r$ is the existence of the so called *mediator*, i.e., a state s such that $q \preceq_F s \preceq_B^{-1} r$ (cf. Fig. 2(a)). The state s can be reached in the same way and in the same context[‡] as r, and, at the same time, the automaton can continue from s in the same way as from q. Hence, intuitively, the newly allowed run based on the jump from r to q does not add anything to the language because it can anyway be realized through s without jumps.

Unfortunately, the relation $\preceq_F \circ \preceq_B^{-1}$ cannot be directly used as it is not transitive, and taking its symmetric closure would thus not yield an equivalence. We thus have to take some of its *transitive fragments*. This is natural as if the automaton can safely jump from q_1 to q_2 and from q_2 to q_3, it should be able to safely jump from q_1 to q_3 too.

This is, however, still not enough. Not all of the transitive fragment of $\preceq_F \circ \preceq_B^{-1}$ can be used for quotienting. We can only take a fragment $\preceq_M$ that is *forward extensible*, meaning

[†]The first view is better when explaining the intuition whereas the other is easier to be used in proofs.

[‡]If a state s is a leaf of a partial run, then by a *context* of s we mean all the other leaves of the partial run.

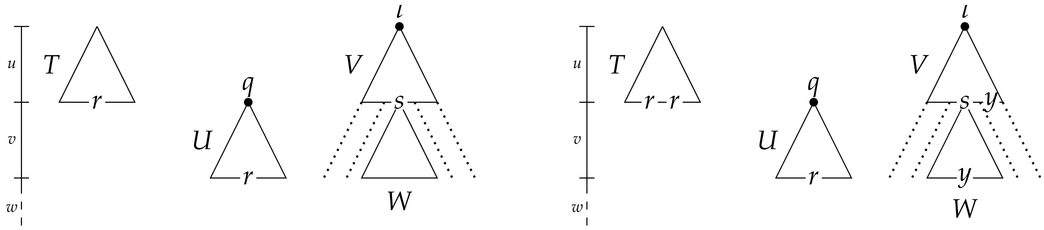

(a) The Mediator (b) Potential Problems

Figure 2: Basic Intuition Behind Mediated Equivalence

that if $q_1 \preceq_M q_2 \preceq_F q_3$, then $q_1 \preceq_M q_3$. The intuitive meaning of this requirement is the following. When a run jumps from r to q, it may be the case that r is again reached later on or it appears in the context of itself (cf. Fig. 2(b)). If r is reached in the continuation of the run from q, the mediated preorder assures that there is some state y in the run continuing from the mediator s that forward simulates r. Similarly, if the context of r contains another occurrence of r, there is some state y in the context of s that forward simulates r. However, this forward simulation is in general guaranteed to hold only when no further jumps are allowed. In order to guarantee a possibility of further simulation, we require that if the computation is allowed to jump from r to q, it is allowed to jump from y to q too.

Finally, to make the mediated equivalence applicable, we must pose one more requirement. Namely, we require that the transitions of the given ABA are not $\preceq_F$-*ambiguous*, meaning that no two states on the right hand side of a transition are forward equivalent. Intuitively, allowing such transitions goes against the spirit of the backward simulation. For a mediator p to backward simulate a state r wrt. rules $\rho_1 : p' \xrightarrow{a} P \cup \{p\}, p \notin P$, and $\rho_2 : r' \xrightarrow{a} R \cup \{r\}, r \notin R$, it must be the case that each state x in the context P of p within ρ_1 is less restrictive (i.e., forward bigger) than some state y in the context R of r within ρ_2. The state r itself is not taken into account when looking for y because we aim at extending its behaviour by collapsing (and it could then become less restrictive than the appropriate x). In the case of $\preceq_F$-ambiguity, the spirit of this restriction is in a sense broken since the forward behaviour of r may still be taken into account when checking that the context of p is less restrictive than that of r. This is because the behaviour of r appears in R as the behaviour of some other state r'' too. Consequently, r and r'' may back up each other in a circular way when checking the restrictiveness of the contexts within the construction of the backward simulation. Both of them can then seem extensible, but once their behaviour gets extended, the restriction of their context based on their own original behaviour is lost, which may then increase the language (an example of such a scenario is given in Appendix F.2). However, in Section 5, we show that $\preceq_F$-ambiguity can be efficiently removed.

Mediated Preorder and Equivalence. Let $\preceq_F$ be a reflexive and transitive forward simulation on $\mathcal{A}$, and $\preceq_B$ a reflexive and transitive backward simulation on $\mathcal{A}$ parameterized by $\preceq_F$. A preorder $\preceq_M \subseteq \preceq_F \circ \preceq_B^{-1}$ such that for all $q, r, s \in Q$, $q \preceq_M r \preceq_F s$ implies $q \preceq_M s$, is a *mediated preorder* induced by $\preceq_F$ and $\preceq_B$. The relation $\equiv_M = \preceq_M \cap \preceq_M^{-1}$ is then a *mediated equivalence* induced by $\preceq_F$ and $\preceq_B$.

LEMMA 4.[3] *There is a unique maximal mediated preorder $\preceq_M$ induced by $\preceq_F$ and $\preceq_B$.*

4.2 Extending Automata According to Mediated Preorder Preserves Language

Quotient Automata versus Extended Automata. We first show that quotienting can be seen as a simpler operation of adding transitions and accepting states. Let $\mathcal{A} = (\Sigma, Q, \iota, \delta, \alpha)$ be an ABA and let $\equiv$ be an equivalence on Q such that $\equiv \; = \; \preceq \cap \preceq^{-1}$ for some preorder $\preceq$. Let the automaton $\mathcal{A}/\equiv$ be the quotient of $\mathcal{A}$ wrt. $\equiv$ that arises by merging $\equiv$-equivalent states of $\mathcal{A}$, and let $\mathcal{A}^+$ be the automaton extended according to $\preceq$, that is created as follows: for every two states q, r of $\mathcal{A}$ with $q \preceq r$, (i) add all outgoing transitions of q to r, (ii) if $q \equiv r$ and q is final, make r final.

The automata $\mathcal{A}/\equiv$ and $\mathcal{A}^+$ are formally defined as follows. Let $Q/\equiv$ denote the quotient of Q wrt. $\equiv$, and let $[q]$ denote the equivalence class of $\equiv$ containing q. Then $\mathcal{A}/\equiv \; = (\Sigma, Q/\equiv, [\iota], \delta/\equiv, \{[q] \mid q \in \alpha\})$ and $\mathcal{A}^+ = (\Sigma, Q, \delta^+, \iota, \alpha^+)$, where $\alpha^+ = \{p \mid \exists q \in \alpha.\, q \equiv p\}$ and, for each $a \in \Sigma$, $q \in Q$, $\delta/\equiv([q], a) = \bigcup_{p \in [q]} \{\{[p'] \mid p' \in P\} \mid P \in \delta(p, a)\}$ and $\delta^+(q, a) = \bigcup_{p \in Q \wedge p \preceq q} \delta(p, a)$. It is not difficult to show that $\mathcal{L}(\mathcal{A}/\equiv) \subseteq \mathcal{L}(\mathcal{A}^+)$ [2] (Lemma 12 in the appendix). Hence, if adding transitions and accepting states according to $\preceq$ preserves the language, then quotienting according to $\equiv$ preserves the language too.

Language Preservation by Mediated Equivalence. We now give a sketch of the proof that extending automata according to the mediated preorder preserves the language. The full proofs can be found in [2] (or Appendix D). For the rest of the section, we fix an ABA $A = (\Sigma, Q, \iota, \delta, \alpha)$, a reflexive and transitive forward simulation $\preceq_F$ on $\mathcal{A}$ such that $\mathcal{A}$ is $\preceq_F$-unambiguous, and a reflexive and transitive backward simulation $\preceq_B$ on $\mathcal{A}$ parameterized by $\preceq_F$. Let $\preceq_M$ be a mediated preorder induced by $\preceq_F$ and $\preceq_B$, and let $\mathcal{A}^+$ be the automaton extended according to $\preceq_M$. Let $\equiv_M \; = \; \preceq_M \cap \preceq_M^{-1}$.

We want to prove that $\mathcal{L}(\mathcal{A}^+) = \mathcal{L}(\mathcal{A})$. The nontrivial part is showing that $\mathcal{L}(\mathcal{A}^+) \subseteq \mathcal{L}(\mathcal{A})$—the converse is obvious. To prove $\mathcal{L}(\mathcal{A}^+) \subseteq \mathcal{L}(\mathcal{A})$, we need to show that, for every accepting run of $\mathcal{A}^+$ on a word w, there is an accepting run of $\mathcal{A}$ on w. We proceed as follows. We first prove Lemma 5, which shows how partial runs of $\mathcal{A}$ with an increased power of their leaves (wrt. $\preceq_F$) can be built incrementally from other runs of $\mathcal{A}$, bridging the gap between $\mathcal{A}$ and $\mathcal{A}^+$. Then we prove Lemma 7 saying that, for every partial run on a word w of $\mathcal{A}^+$, there is a partial run of $\mathcal{A}$ on w that is more accepting (recall that partial runs are finite). By carry this result over to infinite runs we get the proof of Theorem 8.

Consider a partial run T of $\mathcal{A}$ on a word w, we choose for each leaf p of T an $\preceq_M$-smaller state p'. Suppose that we allow p to make one step using the transitions of p' or to become accepting if p' is accepting and $p' \equiv_M p$. (Thus, we give the leaves of T a part of the power they would have in $\mathcal{A}^+$). We will show that there exists a partial run U of $\mathcal{A}$ on w such that (1) it is more accepting than T, and (2) the leaves of U can mimic the next step of the leaves of T even if the leaves of T use their extended power.

The above is formalized in Lemma 5 using the following notation. For a partial run T of $\mathcal{A}$ on w, we define *ext* as an *extension function* that assigns to every branch π of T a state $ext(\pi)$ such that $ext(\pi) \preceq_M leaf(\pi)$.

Let U be a partial run of $\mathcal{A}$ on w. For two branches $\pi \in branches(T)$ and $\psi \in branches(U)$, we say that ψ *strongly covers* π wrt. *ext*, denoted $\pi \preceq_{ext} \psi$, iff $\pi \preceq_\alpha \psi$ and $ext(\pi) \preceq_F leaf(\psi)$. Similarly, we say that ψ *weakly covers* π wrt. *ext*, denoted $\pi \preceq_{\text{w-}ext} \psi$, iff $\pi \preceq_\alpha \psi$ and $ext(\pi) \preceq_M leaf(\psi)$. We extend the concept of covering to partial runs as follows. We write $T \preceq_{ext} U$ (U *strongly covers* T wrt. *ext*) iff $branches(T) \preceq_{ext}^{\forall\exists} branches(U)$ and $root(T) \preceq_B$

root(U). Likewise, we write $T \preceq_{\text{w-}ext} U$ (U *weakly covers* T wrt. *ext*) iff *branches*$(T) \preceq^{\forall\exists}_{\text{w-}ext}$ *branches*(U) and *root*$(T) \preceq_B$ *root*(U). Note that we have $\preceq_{ext} \subseteq \preceq_{\text{w-}ext}$ for branches as well for partial runs because $\preceq_F \subseteq \preceq_M$. So, the strong covering implies the weak one.

LEMMA 5. *For any partial run T of $\mathcal{A}$ on a word w with an extension function ext, there is a partial run U of $\mathcal{A}$ on w with $T \preceq_{ext} U$.*

Proving Lemma 5 is the most intricate part of the proof of Theorem 8. We introduce the concepts used within the proof of Lemma 5 and provide an overview of the proof.

If $T \preceq_{ext} T$, we are done as in the statement of the lemma, we can take T to be U. So, suppose that $T \not\preceq_{ext} T$. Observe that *root*$(T) \preceq_B$ *root*(T), and every branch of T weakly covers itself, which means that $T \preceq_{\text{w-}ext} T$. We will show how to reach U by a chain of partial runs derived from T. The partial runs within the chain will all weakly cover T. Runs further from T will in some sense cover T more strongly than the runs closer to T. The last partial run of the chain will cover T strongly. To do this, we need a suitable measure that, for a partial run V of $\mathcal{A}$ on w with $T \preceq_{\text{w-}ext} V$, tells us how strongly V covers T.

To define the measure, we concentrate on branches of V that cause that V does not cover T strongly. These are branches $\psi \in$ *branches*(V) for which there is no $\pi \in$ *branches*(T) with $\pi \preceq_{ext} \psi$ (there are only some $\pi \in$ *branches*(T) with $\pi \preceq_{\text{w-}ext} \psi$). We call them *strict weakly covering branches*. Let $\mathsf{sw}_T(V)$ denote the tree which is the subset of V containing prefixes of strict weakly covering branches of V wrt. T. Note that $T \preceq_{ext} V$ iff V contains no strict weakly covering branches, which is equivalent to $\mathsf{sw}_T(V) = \emptyset$. For a partial run W of $\mathcal{A}$ on w, we will define which of V and W cover T more strongly by comparing $\mathsf{sw}_T(V)$ and $\mathsf{sw}_T(W)$. For this, we need the following definitions.

Given a finite tree X over Q and $\tau \in Q^+$, we define the *tree decomposition* of X according to τ as the sequence of (finite) sets of paths $\langle \tau, X\rangle = X \ominus_1 \tau, X \ominus_2 \tau, \ldots, X \ominus_{|\tau|} \tau$. We also let $\langle \epsilon, X\rangle =$ *branches*(X), which is a sequence of length 1. Notice that under the condition that $\tau \notin$ *branches*(X), $\langle \tau, X\rangle = \emptyset \ldots \emptyset$ implies that $X = \emptyset$[§].

Let $\tau_V \in V \cup \{\epsilon\}$ and $\tau_W \in W \cup \{\epsilon\}$ be such that $\tau_V \notin$ *branches*$(\mathsf{sw}_T(V))$ and $\tau_W \notin$ *branches*$(\mathsf{sw}_T(W))$. We say that W *covers T more strongly* than V wrt. τ_V and τ_W, denoted $V \prec^T_{\tau_V,\tau_W} W$, iff *root*$(V) \preceq_B$ *root*(W) and $\langle \tau_V, \mathsf{sw}_T(V)\rangle \sqsubset \langle \tau_W, \mathsf{sw}_T(W)\rangle$, where $\sqsubset$ is a binary relation on sequences of sets of paths defined as follows.

For two sets of paths P and P', we use $P \prec^{\forall\exists}_F P'$ to denote that $P \preceq^{\forall\exists}_F P'$ but not $P' \preceq^{\forall\exists}_F P$. In other words, the upward closure of P' wrt. $\preceq_F$ is a proper subset of the upward closure of P wrt. $\preceq_F$. Then, for sequences of finite sets $S, S' \in (2^Q)^+$, $S \sqsubset S'$ iff there is some $k \in \mathbb{N}, k \leq \min\{|S|, |S'|\}$, such that $S_k \prec^{\forall\exists}_F S'_k$ and for all $1 \leq j < k$, $S_j \preceq^{\forall\exists}_F S'_j$. It is not hard to show that the relation $\sqsubset$ is a partial order. Observe that $\sqsubset$ does not allow infinite increasing chains of sequences where the length of the sequences is bounded by some constant (this follows from that $\preceq_F$ compares only paths of an equal length and therefore every increasing chain of finite sets of paths related by $\prec^{\forall\exists}_F$ is finite). Moreover, $S \sqsubset \emptyset \ldots \emptyset$ for every sequence of sets of paths $S \neq \emptyset \ldots \emptyset$.

[§]Note that if $\tau \in$ *branches*(X), $\langle \tau, X\rangle = \emptyset \ldots \emptyset$ does not imply $X = \emptyset$ as τ could be the only branch of X. This is important as for a partial run Y and $\tau' \in Y$, if $\tau' \notin$ *branches*(Y), the implications $\langle \tau', \mathsf{sw}_T(Y)\rangle = \emptyset \ldots \emptyset \implies \mathsf{sw}_T(Y) = \emptyset \implies T \preceq_{ext} Y$ hold. However, the first implication does not hold if $\tau' \in$ *branches*(Y).

LEMMA 6. *Given a partial run V of $\mathcal{A}$ on w s.t. $T \preceq_{\text{w-}ext} V$, $T \not\preceq_{ext} V$, and $\tau_V \in V \cup \{\epsilon\}$ with $\tau_V \notin branches(\mathsf{sw}_T(V))$, we can construct a partial run W of $\mathcal{A}$ on w with $T \preceq_{\text{w-}ext} W$ and a path $\tau_W \in W$ with $\tau_W \notin branches(\mathsf{sw}_T(W))$ such that $V \prec^T_{\tau_V,\tau_W} W$.*

PROOF. [Sketch] The proof of Lemma 6 relies on Lemma 3 and the definition of $\preceq_M$. We first choose a suitable branch π of $\mathsf{sw}_T(V)$ as follows. Let $1 \leq k \leq |\tau_V|$ be some index such that $\mathsf{sw}_T(V) \ominus_k \tau_V$ is nonempty. If $\tau_V = \epsilon$, then $k = 1$. We choose some $\pi' \in \mathsf{sw}_T(V) \ominus_k \tau_V$ which is minimal wrt. $\preceq_F$, meaning that there is no $\pi'' \in \mathsf{sw}_T(V) \ominus_k \tau_V$ different from π' such that $\pi'' \preceq_F \pi'$. We put $\pi = \tau_V^k \pi'$. We note that this is the place where we use the $\preceq_F$-unambiguity assumption. If $\mathcal{A}$ was $\preceq_F$-ambiguous, there need not be a k such that $\mathsf{sw}_T(V) \ominus_k \tau_V$ contains a minimal element wrt. $\preceq_F$.

From $ext(\pi) \preceq_M leaf(\pi)$, there is a mediator s with $ext(\pi) \preceq_F s \succeq_B leaf(\pi)$. We apply Lemma 3 to V, π, $leaf(\pi)$ and s, which give us a partial run W and $\psi \in branches(W)$ with $leaf(\psi) = s$ such that $\pi \preceq_B \psi$, and for all $1 \leq i \leq |\pi|$, $V \ominus_i \pi \preceq_F^{\forall\exists} W \ominus_i \psi$. Let $\tau_W = \psi$. The proof can be concluded by showing that (i) $T \preceq_{\text{w-}ext} W$, (ii) $\tau_W \notin branches(\mathsf{sw}_T(W))$, and (iii) $\langle \tau_V, \mathsf{sw}_T(V)\rangle \sqsubset \langle \tau_W, \mathsf{sw}_T(W)\rangle$, which implies $V \prec^T_{\tau_V,\tau_W} W$. ∎

Now we construct a run U strongly covering T as follows. Starting from T and ϵ, we can construct a chain $T \prec^T_{\epsilon,\tau_1} T_1 \prec^T_{\tau_1,\tau_2} T_2 \prec^T_{\tau_2,\tau_3} T_3 \ldots$ by successively applying Lemma 6 for each i, $\tau_i \in T_i$, $\tau_i \notin branches(\mathsf{sw}_T(T_i))$, and $T \preceq_{\text{w-}ext} T_i$. Observe that by the definition of stronger covering, we have that $\langle \epsilon, \mathsf{sw}_T(T)\rangle \sqsubset \langle \tau_1, \mathsf{sw}_T(T_1)\rangle \sqsubset \langle \tau_2, \mathsf{sw}_T(T_2)\rangle \sqsubset \langle \tau_3, \mathsf{sw}_T(T_3)\rangle \ldots$ Notice that, for each i, as $T \preceq_{\text{w-}ext} T_i$, $height(T_i) = height(T)$. Therefore the length of τ_i as well as the length of $\langle \tau_i, \mathsf{sw}_T(T_i)\rangle$ are bounded by $height(T)$.

Recall that (i) the relation $\sqsubset$ is a partial order, (ii) that $\sqsubset$ does not allow infinite increasing chains of sequences where the length of the sequences is bounded by some constant, and (iii) that $S \sqsubset \emptyset \ldots \emptyset$ for every sequence $S \neq \emptyset \ldots \emptyset$. This means that after a finite number of steps, this chain must arrive to its last T_k and τ_k with $\langle \tau_k, \mathsf{sw}_T(T_k)\rangle = \emptyset \ldots \emptyset$. This means that $\mathsf{sw}_T(T_k) = \emptyset$, which implies that $T \preceq_{ext} T_k$. We can put $U = T_k$ and Lemma 5 is proven.

Now we can use Lemma 5 to prove Lemma 7. It relates partial runs of $\mathcal{A}^+$ with partial runs of $\mathcal{A}$ by the relation $\preceq_{\alpha^+ \Rightarrow \alpha}$ defined as follows. For two states q and r, $q \preceq_{\alpha^+ \Rightarrow \alpha} r$ iff $q \in \alpha^+ \implies r \in \alpha$. For two paths $\pi, \psi \in Q^+$, $\pi \preceq_{\alpha^+ \Rightarrow \alpha} \psi$ iff $|\pi| = |\psi|$ and for all $1 \leq i \leq |\pi|$, $\pi_i \in \alpha^+ \implies \psi_i \in \alpha$. Finally, for finite trees T and U over Q, we use $T \preceq_{\alpha^+ \Rightarrow \alpha} U$ to denote that $branches(T) \preceq^{\forall\exists}_{\alpha^+ \Rightarrow \alpha} branches(U)$.

LEMMA 7. *For any partial run T of $\mathcal{A}^+$ on $w \in \Sigma^\omega$, there exists a partial run U of $\mathcal{A}$ on w such that $root(T) \preceq_B root(U)$ and $T \preceq_{\alpha^+ \Rightarrow \alpha} U$.*

The proof of Lemma 7 is done by induction on the structure of T, where the induction step employs Lemma 5 (which bridges the gap between $\mathcal{A}^+$ and $\mathcal{A}$ by showing that there is a partial run of $\mathcal{A}$ strongly covering T even when the power of its leaves is extended by transitions of some $\preceq_M$-smaller states). With Lemma 7 in hand, we can prove that for each accepting run of $\mathcal{A}^+$ on a word w, there is an accepting run of $\mathcal{A}$ on w. This requires to carry Lemma 7 from finite partial runs to full infinite runs[¶]. This results in Theorem 8, which together with the fact that $\mathcal{L}(\mathcal{A}/\equiv) \subseteq \mathcal{L}(\mathcal{A}^+)$ immediately gives Corollary 9.

[¶] For an accepting run T of $\mathcal{A}^+$ on a word w, Lemma 7 gives us for every $k \in \mathbb{N}$ and a prefix of T of the height k a partial run of U of the same height that is more accepting. From the infinite set of partial runs of $\mathcal{A}$ obtained this way, we can construct an accepting run of $\mathcal{A}$ on w. The details may be found in [2] and in Appendix D.3.

THEOREM 8. $\mathcal{L}(\mathcal{A}^+) = \mathcal{L}(\mathcal{A})$.

COROLLARY 9. *Quotienting with mediated equivalence preserves the language.*

5 Algorithm for Computing Mediated Preorder

In this section, we describe an algorithm for computing mediated preorder on an ABA $\mathcal{A} = (\Sigma, Q, \iota, \delta, \alpha)$. We first explain how to compute the maximal forward simulation $\preceq_F$ and backward simulation $\preceq_B$ of $\mathcal{A}$. Both $\preceq_F$ and $\preceq_B$ will be used as the input parameters for computing the mediated preorder $\preceq_M$. In the rest of the section, we will fix $\mathcal{A}$ as the input ABA, use n for the number of states in $\mathcal{A}$, and use m for the number of transitions in $\mathcal{A}$.

Forward Simulation. The algorithm for computing maximal forward simulation $\preceq_F$ on $\mathcal{A}$ can be found in Fritz and Wilke's work [5] (it is called direct simulation in their paper). They reduce the problem of computing maximal forward simulation to a simulation game. Although Fritz and Wilke use a slightly different definition of ABA, it is easy to translate $\mathcal{A}$ to an ABA under their definition with $O(n+m)$ states and $O(nm)$ transitions and then use their algorithm to compute $\preceq_F$. The time complexity of the above procedure is $O(nm^2)$.

Removing Ambiguity. As shown in Section 4.1, $\mathcal{A}$ needs to be $\preceq_F$-unambiguous for mediated minimization. Here we describe how to modify $\mathcal{A}$ to make it not $\preceq_F$-ambiguous. The modification does not change the the language of $\mathcal{A}$ and also the forward simulation relation $\preceq_F$, therefore we do not need to recompute forward simulation again for the modified automaton.

Here we describe the ambiguity removal procedure. For every transition $p \xrightarrow{a} P$ with $P = \{p_1, \ldots, p_k\}$ and for each $i \in \{1, \ldots, k\}$, we check if there exists some $i < j \leq k$ such that $p_j \preceq_F p_i$. If there is one, remove p_i from P. This procedure has time complexity $O(n^2 m)$.

Backward Simulation. We now show how to translate the problem of computing maximal backward simulation to a problem of computing maximal simulation on a labeled transition system.

Computing Simulation on Labeled Transition Systems. Let $T = (S, \mathcal{L}, \rightarrow)$ be a finite *labeled transition system (LTS)*, where S is a finite set of states, $\mathcal{L}$ is a finite set of labels, and $\rightarrow \subseteq S \times \mathcal{L} \times S$ is a transition relation. A *simulation* on T is a binary relation $\preceq_L$ on S such that if $q \preceq_L r$ and $(q, a, q') \in \rightarrow$, then there is an r' with $(r, a, r') \in \rightarrow$ and $q' \preceq_L r'$.

Here we describe the problem of computing the maximal simulation on an LTS. Given an LTS $T = (S, \mathcal{L}, \rightarrow)$ and an *initial* preorder $I \subseteq S \times S$, the task is to find out the unique maximal simulation on T included in I. An algorithm for computing maximal simulation $\preceq^I$ on the LTS T included in I with time complexity $O(|\mathcal{L}|.|S|^2 + |S|.|\rightarrow|)$ and space complexity $O(|\mathcal{L}|.|S|^2)$ can be found in [1].

Computing Backward Simulation via a Reduction to LTS. The problem of computing the maximal backward simulation on $\mathcal{A}$ can be reduced to the problem of computing simulation on an LTS. In order to simplify the explanation of the reduction, we first make the following definition. An *environment* is a tuple of the form $(p, a, P \setminus \{p'\})$ obtained by removing a state $p' \in P$ from the transition $p \xrightarrow{a} P$ of $\mathcal{A}$. Intuitively, an environment records the neighbors of the removed state p' in the transition $p \xrightarrow{a} P$. We denote the set of all environments of $\mathcal{A}$ by $Env(A)$. Formally, we define the LTS $A^{\odot} = (Q^{\odot}, \Sigma, \Delta^{\odot})$ as follows:

$$\begin{array}{ccccccc} \text{A transition in } \mathcal{A} & & & & \text{Transitions in } \mathcal{A}^{\odot} & & \\ & & p_1^{\odot} & \xrightarrow{a} & (p,a,\{p_2,p_3\})^{\odot} & \searrow^{a} & \\ p \xrightarrow{a} \{p_1,p_2,p_3\} & \Rightarrow & p_2^{\odot} & \xrightarrow{a} & (p,a,\{p_1,p_3\})^{\odot} & \xrightarrow{a} & p^{\odot} \\ & & p_3^{\odot} & \xrightarrow{a} & (p,a,\{p_1,p_2\})^{\odot} & \nearrow_{a} & \end{array}$$

Figure 3: An example of the reduction from an ABA transition to LTS transitions

- $Q^{\odot} = \{q^{\odot} \mid q \in Q\} \cup \{(p,a,P)^{\odot} \mid (p,a,P) \in Env(A)\}$.
- $\Delta^{\odot} = \{(p,a,P \setminus \{p'\})^{\odot} \xrightarrow{a} p^{\odot}, p'^{\odot} \xrightarrow{a} (p,a,P \setminus \{p'\})^{\odot} \mid P \in \delta(p,a), p' \in P\}$.

An example of the reduction is given in Figure 3. The goal of this reduction is to obtain a simulation relation on $A^{\odot}$ with the following property: $p^{\odot}$ is simulated by $q^{\odot}$ in $A^{\odot}$ iff $p \preceq_B q$ in $\mathcal{A}$. However, the maximal simulation on $A^{\odot}$ is not sufficient to achieve this goal. Some essential conditions for backward simulation (e.g., $p \preceq_B q \implies p \preceq_\alpha q$) are missing in $A^{\odot}$. This can be fixed by defining a proper initial preorder I.

Formally, we define $I = \{(q_1^{\odot}, q_2^{\odot}) \mid q_1 \preceq_\iota q_2 \wedge q_1 \preceq_\alpha q_2\} \cup \{((p,a,P)^{\odot},(r,a,R)^{\odot}) \mid P \preceq_F^{\forall\exists} R\}$. Observe that I is a preorder. Recall that according to the definition of the backward simulation, $p \preceq_B r$ implies that (1) $p \preceq_\iota r$, (2) $p \preceq_\alpha r$, and (3) for all transitions $q \xrightarrow{a} P \cup \{p\}, p \notin P$, there exists a transition $s \xrightarrow{a} R \cup \{r\}, r \notin R$ such that $q \preceq_B s$ and $P \preceq_F^{\forall\exists} R$. The set $\{(q_1^{\odot}, q_2^{\odot}) \mid q_1 \preceq_\iota q_2 \wedge q_1 \preceq_\alpha q_2\}$ encodes the conditions (1) and (2) required by the backward simulation, while the set $\{((p,a,P)^{\odot},(r,a,R)^{\odot}) \mid P \preceq_F^{\forall\exists} R\}$ encodes the condition (3). A simulation relation $\preceq^I$ can be computed using the aforementioned procedure with LTS $A^{\odot}$ and the *initial* preorder I. The following theorems shows the correctness and complexity of computing backward simulation.

THEOREM 10. *For all $q, r \in Q$, we have $q \preceq_B r$ iff $q^{\odot} \preceq^I r^{\odot}$.*

THEOREM 11. *Computing backward simulation has both time and space complexity $O(nm^3)$.*

The complexity comes from three parts of the procedure: (1) compiling $\mathcal{A}$ into its corresponding LTS $A^{\odot}$, (2) computing the initial preorder I, and (3) running the algorithm for computing the LTS simulation relation. The LTS $A^{\odot}$ has at most nm+n states and $2nm$ transitions. It follows that Part (3) has time complexity $O(|\Sigma|n^2m^2)$ and space complexity $O(|\Sigma|n^2m^2)$. In Appendix E, we show that among the three parts, Part (3) has the highest time[‖] and space complexity and therefore computing backward simulation also has time complexity $O(|\Sigma|n^2m^2)$ and space complexity $O(|\Sigma|n^2m^2)$. Under our definition of ABA, every state has at least one outgoing transition for each symbol in Σ. It follows that $m \geq |\Sigma|n$. Therefore, we can also say that the procedure for computing maximal backward simulation has time complexity $O(nm^3)$ and space complexity $O(nm^3)$.

Mediated Preorder. Here we explain how to compute the mediated preorder $\preceq_M$ of $\mathcal{A}$ from $\preceq_F$ and $\preceq_B$. It is proved in [1] that $\preceq_M$ equals the maximal relation $R \subseteq \preceq_F \circ \preceq_B^{-1}$ satisfying $x\ R\ y \preceq_F z \implies x\ (\preceq_F \circ \preceq_B^{-1})\ z$. Based on the result, we can obtain the mediated preorder by the following procedure. Initially, let $\preceq_M = \preceq_F \circ \preceq_B^{-1}$. For all $(p,q) \in \preceq_M$, if there exists some $(q,r) \in \preceq_F$ such that $(p,r) \notin \preceq_F \circ \preceq_B^{-1}$, remove (p,q) from $\preceq_M$. A naive implementation of this simple procedure has time complexity $O(n^3)$.

[‖]In [2] we will describe an efficient algorithm for computing I. It has time complexity $O(n^2m^2)$ and space complexity $O(n^2m^2)$.

6 Experimental Results

In this section, we evaluate the performance of mediated minimization by applying it to accelerate the algorithm proposed by Vardi and Kupferman [9] for complementing nondeterministic Büchi automata (NBA). In this algorithm, ABA's are used as intermediate notion for the complementation. To be more specific, the complementation algorithm has two steps: (1) it translates an NBA to an ABA that recognizes its complement language, and (2) it translates the ABA back to an equivalent NBA. The second step is an exponential procedure (exponential in the size of the ABA), hence reducing the size of the ABA before the second step usually pays off.

The experimentation is carried out as follows. Three sets of 100 random NBA's (of $|\Sigma| =$ 2,4, and 8, respectively) are generated by the GOAL [11] tool and then used as inputs of the complementation experiments. We compare results of experiments performed according to the following different options: (1) **Original:** keep the ABA as what it is, (2) **Mediated:** minimizing the ABA with mediated equivalence, and (3) **Forward:** minimizing the ABA with forward equivalence.

For each input NBA, we first translate it to an ABA that recognizes its complement language. The ABA is (1) processed according to one of the options described above and then (2) translated back to an equivalent NBA using an exponential procedure **. The results are given in Table 1 and Table 2. Table 1 is an overall comparison between the three different options and Table 2 is a more detailed comparison between **Mediated** and **Forward** minimization.

In Table 1, the columns "NBA" and "Complemented-NBA" are the average statistical data of the input NBA and the complemented NBA. The column "Time(ms)" is the average execution time in milliseconds. "Timeout" is the number of cases that cannot finish within the timeout period (10 min). Note that in the table, the cases that cannot finish within the timeout period are excluded from the average number. From this table, we can see that minimization by mediated equivalence can effectively speed up the complementation and also reduce the size of the complemented NBA's.

	$\|\Sigma\|$	NBA		Complemented-NBA		Time (ms)	Timeout (10 min)
		St.	Tr.	St.	Tr.		
Original	2	2.5	3.3	13.9	52.75	5500.9	0
Mediated				6.68	34.02	524.7	0
Forward				9.45	55.25	5443.7	1
Original	4	3.3	6.0	46.4	348.5	9298.6	6
Mediated				20.42	235.5	1985.4	6
Forward				26.88	325.6	1900.6	7
Original	8	4.7	11.9	127.1.3	1723.4	33429.4	24
Mediated				57.63	1738.3	12930.6	21
Forward				81.23	2349.2	22734.2	24

Table 1: Combining minimization with complementation.

In Table 2, we compare the performance between **Mediated** and **Forward** minimization in detail. The columns "Minimized-ABA" and "Complemented-NBA" are the average difference in the sizes of the ABA after minimization and the complemented BA. From the table, we observe that mediated minimization results in a much better reduction than forward minimization.

	$\|\Sigma\|$	Minimized-ABA		Complemented-NBA	
		St.	Tr.	St.	Tr.
Average Difference	2	33.54%	51.62%	63.3%	235.56%
	4	36.24%	51.44%	89.9%	298.99%
	8	27.94%	40.88%	152.3%	412.7%

Table 2: Comparison: **Mediated** v.s. **Forward**

**For the option "Original", we also use the optimization suggested in [9] that only takes consistent subset.

7 Conclusion and Future Work

We combined forward and backward simulation to form a coarser relation called mediated preorder and showed that quotienting wrt. mediated equivalence preserves the language of ABA. Moreover, we developed an efficient algorithm for computing mediated equivalence. Experimental results show that the mediated reduction of ABA significantly outperforms the reduction based on forward simulation.

In the future, we would like to extend our experiments to other applications such as LTL to NBA translation. Furthermore, we would like to extend the mediated equivalence by building it on top of even coarser forward simulation relations, e.g., *delayed* or *fair* forward simulation relations [6]. Also, we would like to study the possibility of using mediated preorder to remove redundant transitions (similar to the approaches described in [10]). We believe that the extensions described above can significantly improve the performance of mediated reduction.

Acknowledgements. We thank the anonymous reviewers for their useful comments. The work was supported by the CONNECT project, the UPMARC project, the Czech Science Foundation (projects 102/07/0322, 102/09/H042), the Barrande project MEB 020840, and the Czech institutional project MSM 0021630528.

References

[1] P. Abdulla, A. Bouajjani, L. Holík, L. Kaati, and T. Vojnar. Computing Simulations over Tree Automata: Efficient Techniques for Reducing Tree Automata. In *Proc. of TACAS'08*, LNCS 4963, pages 93–108. Springer, 2008.

[2] P. Abdulla, Y.-F. Chen, L. Holík, and T. Vojnar. Mediating for reduction (on minimizing alternating büchi automata). Tech. Rep. FIT-TR-2009-02, Brno University of Technology, 2009.

[3] P. Abdulla, L. Holík, L. Kaati, and T. Vojnar. A uniform (bi-)simulation-based framework for reducing tree automata. *ENTCS*, 2009(251):27–48, 2009.

[4] A. Farzan, Y.-F. Chen, E. Clarke, Y.-K. Tsay, and B.-Y. Wang. Extending automated compositional verification to the full class of omega-regular languages. In *Proc. of TACAS'08*, LNCS 4963, pages 2–17. Springer, 2008.

[5] C. Fritz and T. Wilke. State space reductions for alternating Büchi automata: Quotienting by simulation equivalences. In *Proc. of FSTTCS'02*, LNCS 2556, pages 157–169. Springer, 2002.

[6] C. Fritz and T. Wilke. Simulation relations for alternating Büchi automata. *Theoretical Computer Science*, 338(1-3):275–314, 2005.

[7] P. Gastin and D. Oddoux. Fast LTL to Büchi automata translations. In *Proc. of CAV'01*, LNCS 2102, pages 53–65. Springer, 2001.

[8] S. Gurumurthy, O. Kupferman, F. Somenzi, and M. Vardi. On complementing nondeterministic Büchi automata. In *Proc. of CHARME'03*, LNCS 2860, pages 96–110. Springer, 2003.

[9] O. Kupferman and M. Vardi. Weak alternating automata are not that weak. *ACM Transactions on Computational Logic*, 2(3):408–429, 2001.

[10] F. Somenzi and R. Bloem. Efficient Büchi automata from LTL formulae. In *Proc. of CAV'00*, LNCS 1855. Springer, 2000.

[11] Y.-K. Tsay, Y.-F. Chen, M.-H. Tsai, K.-N. Wu, and W.-C. Chan. GOAL: A graphical tool for manipulating Büchi automata and temporal formulae. In *Proc. of TACAS'07*, LNCS 4424, pages 466–471. Springer, 2007.

[12] M. Vardi. Automata-theoretic model checking revisited. In *Proc. of VMCAI'07*, LNCS 4349, pages 137–150. Springer, 2007.

A Initial Preorder for Computing Backward Simulation

As mentioned in the main text, we need to compute a proper *initial preorder* I for the reduction from the problem of backward simulation on an ABA $\mathcal{A} = (\Sigma, Q, \iota, \delta, \alpha)$ to problem of simulation on the LTS $A^{\odot} = (Q^{\odot}, \Sigma, \Delta^{\odot})$. The preorder I is the union of two sets: $\{(q_1^{\odot}, q_2^{\odot}) \mid q_1 \preceq_\iota q_2 \wedge q_1 \preceq_\alpha q_2\}$ and $\{((p,a,P)^{\odot}, (r,a,R)^{\odot}) \mid \forall r_j \in R. \exists p_i \in P. p_i \preceq_F r_j\}$. Let n and m be the number of states and transitions in $\mathcal{A}$, respectively. It is trivial that the first set can be computed by an algorithm with time complexity $O(n^2)$. However, a naive algorithm (via a pairwise comparison of all different environments in $env(\mathcal{A})$) for computing the second set has time complexity $O(n^4m^2)$. Here we will describe a more efficient algorithm, which allows the computation of I to have both time and space complexity $O(n^2m^2)$.

The main idea of the algorithm is the following. For each pair of two given transitions, it examines all pairs of related environments at once and adds pairs of states in $A^{\odot}$ to I when needed. This action has both time and space complexity $O(n^2)$. Because $\mathcal{A}$ has at most m^2 different pairs of transitions, the second set of I can be computed by the new algorithm with both time and space complexity $O(n^2m^2)$.

In the rest of this section, we will explain how to efficiently compute all pairs of environments that should be added to I at once from two given transitions. For each pair of transitions $p \xrightarrow{a} P$ and $r \xrightarrow{a} R$, we maintain a mapping function $\beta : R \to \{T, F\} \cup P$ such that

$$\beta(r') = \begin{cases} T & \text{if there exsit more than two states in } P \text{ that are forward smaller than } r'. \\ F & \text{if there exsits no state in } P \text{ that is forward smaller than } r'. \\ p' & \text{if } p' \text{ is the only state in } P \text{ such that } p' \preceq_F r'. \end{cases}$$

The mapping function β can be computed by Algorithm 1 with both time and space complexity $O(n^2)$.

Algorithm 1: *Generate a Mapping Function For Two Transitions*

Input: Two transitions $p \xrightarrow{a} P$ and $r \xrightarrow{a} R$ in $\mathcal{A}$.
Output: A mapping function $\beta : P \to R \cup \{T, F\}$.
`/* initialization */`
forall $r' \in R$ **do** $\beta(r') := F$;
forall $p' \in P, r' \in R$ **do**
  if $p' \preceq_F r'$ **then**
    if $\beta(r') = F$ **then** $\beta(r') := p'$;
    else $\beta(r') := T$;

Let us consider a pair of states $((p, a, P \setminus \{p'\})^{\odot}, (r, a, R \setminus \{r'\})^{\odot})$ in $\mathcal{A}^{\odot}$. This pair can be added to I if and only if the following two conditions hold:

1. $\forall \hat{r} \in (R \setminus \{r'\}). \beta(\hat{r}) \neq F$.
2. $\forall \hat{r} \in (R \setminus \{r'\}). \beta(\hat{r}) \neq p'$.

With some preprocessing, both conditions can be efficiently checked in constant time. Although the preprocessing has time complexity and space complexity $O(n)$, it has to be done only at the beginning of the algorithm and can then be reused to check all pairs of environments generated from the given pair of transitions.

We need the following preprocessing for condition 1. We define $\hat{p} \in P$ as the `KeyState` if $\hat{p}$ is the only one state in P such that $\beta(\hat{p}) = F$. Given a mapping function β, the `KeyState` can be found efficiently (with time complexity $O(n)$ and space complexity $O(1)$) by scanning through R and

- if there exist two states $r_1, r_2 \in R$ such that $\beta(r_1) = \beta(r_2) = F$, the algorithm terminates immediately because it follows that none of the pairs of environments generated from the given pair of transitions satisfies the requirement of I.
- if there exists only one state such that β maps it to F, let it be the `KeyState`.

Then we have condition 1 is satisfied if (1) there is no `KeyState` or (2) r' is the `KeyState`.

The preprocessing for condition 2 is the following. We maintain a mapping function $\gamma : P \to \{T, F\} \cup R$ (similar to the reverse function of β) such that

$$\gamma(p') = \begin{cases} T & \text{if } |\{\hat{r} \mid \hat{r} \in R \wedge \beta(\hat{r}) = p'\}| > 1 \\ F & \text{if } |\{\hat{r} \mid \hat{r} \in R \wedge \beta(\hat{r}) = p'\}| = 0 \\ r' & \text{if } |\{\hat{r} \mid \hat{r} \in R \wedge \beta(\hat{r}) = p'\}| = 1 \wedge \beta(r') = p'. \end{cases}$$

The mapping function γ can be found with both time complexity $O(n)$ and space complexity $O(n^2)$ by scanning through β. With the function γ, condition 2 can be easily verified by checking if $\gamma(p') \in \{F, r'\}$, which means that for all the states $\hat{r}$ in $R \setminus \{r'\}$, p' is not the only state such that $p' \preceq_F \hat{r}$.

Algorithm 2: *Add Pairs of States to I*

Input: Transitions $p \xrightarrow{a} P, r \xrightarrow{a} R$ in $\mathcal{A}$ and the corresponding mapping function β.
`/* Preprocessing for condition 1 */`
forall $r' \in R$ **do** **if** $\beta(r') = F$ **then**
  if *there is no* `KeyState` **then** Let r' be the `KeyState`;
  else Terminate the algorithm;
`/* Preprocessing for condition 2 */`
forall $p' \in P$ **do** $\gamma(p') := F$;
forall $r' \in R$ **do** **if** $\beta(r') \notin \{T, F\}$ **then**
  if $\gamma(\beta(r')) = F$ **then** $\gamma(\beta(r')) := r'$;
  else $\gamma(\beta(r')) := T$;
`/* main loop */`
forall $p' \in P, r' \in R$ **do**
  if *there is no* `KeyState` *or* r' *is the* `KeyState` **then**
    if $\gamma(p') \in \{F, r'\}$ **then** add $((p, a, P \setminus \{p'\})^{\odot}, (r, a, R \setminus \{r'\})^{\odot})$ to I

In Algorithm 2, we first find out the `KeyState` if there is one and compute the function γ from β. Then in the main loop, for each pair of states $((p, a, P \setminus \{p'\})^{\odot}, (r, a, R \setminus \{r'\})^{\odot})$, we check if it can be included to I by verifying the two conditions that we mentioned before. Now it is easy to see that the Algorithm 2 has both time and space complexity $O(n^2)$. It follows that the initial preorder I can be computed with both time and space complexity $O(n^2 m^2)$.

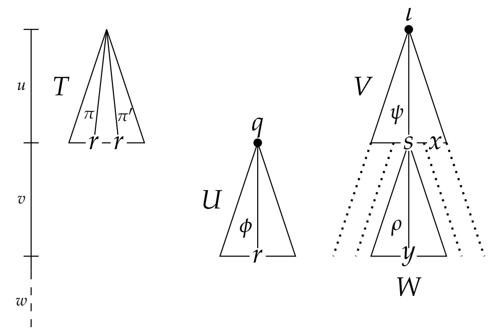

Figure 4: Potential Problems When $\preceq_M$ Is Not Forward Extensible

B Potential Problems When $\preceq_M$ Is Not Forward Extensible

Here we describe in detail the potential problems when $\preceq_M$ is not forward extensible (see Figure B for the illustrations).

Problem (i): The first problem will arise if there is a branch ϕ of U with $leaf(\phi) = r$. Here, apart from interconnecting T and U, r can use its new transitions also at the end of $\pi\phi$ and connect another copy of U to the end of $\pi\phi$. Suppose that all leaves of T except r accept vvw and that all leaves of U except r accept vw. Then this enables a new accepting run on the word $uvvw$. In this case, the existence of the mediator s is not a guarantee that some accepting run on $uvvw$ was possible before adding transitions to r.

Problem (ii): Another problem may arise if there are two (or more) branches in T ending by r. Here we use the two branches π and π' in Figure B as an example. To construct an accepting run on uvw from T, r has to use the transitions of q at the end of π as well as at the end of π' to connect U to T in the both places. But partial run V "covers" only one of the two occurrences of r. There may be a leaf x of V different from s for which r is the only leaf in T with $r \preceq_F x$. Therefore, x needs not to accept vw as there is no guaranteed relation between q and x in which case V is not a prefix of an accepting run on uvw and uvw need not be in $\mathcal{L}(\mathcal{A})$. Note that a very similar situation can arise while attempting to quotient using pure backward simulation equivalence which is the main reason why it cannot be used.

The solution of the both problems is to allow r to use the transitions of q only if $q \preceq r$, where $q \preceq r$ means that (a) there is a mediator for q and r and (b) for any state t, $r \preceq_F t$ implies $q \preceq t$. We will show how the assumption of $q \preceq r$ helps to solve Problem (i) and (ii).

In the case of Problem (i), if y uses transitions of q to accept vw, then W becomes a prefix of an accepting run on vvw and thus V becomes a prefix of a new accepting run on $uvvw$. We know that $r \preceq_F y$. Thus, according to the definition of $\preceq$, $q \preceq r \preceq_F y$ gives $q \preceq y$, which implies that there is a mediator for q and y. Observe that y used transitions of q just once. Therefore, by an analogical argument by which we derived that $\mathcal{A}$ accepts uvw in the first case when r used the new transitions only once, we can here derive that there is an accepting run of $\mathcal{A}$ on $uvvw$ which does not involve new transitions.

In the case of Problem (ii), if x uses the transitions of q to accept vw, V becomes a prefix of a new accepting run on uvw. We know that $r \preceq_F x$ and thus $q \preceq r \preceq_F x$ gives $q \preceq x$, which means that there is a mediator for q and x. Similarly as in the previous case, as x used the transitions of q only once, we can derive that there exists an accepting run of $\mathcal{A}$ on uvw that does not involve new transitions.

The argumentation from the two above paragraphs can be used inductively for a run where r uses transitions of q arbitrarily many times.

C Basic Properties of Simulation Relations

Here we give the proofs lemmas from Section 3.

PROOF. [Lemma 2] We prove the lemma by induction on $height(T)$. In the base case when $T = \{p\}$, it is sufficient to take $U = \{r\}$. Suppose now that the lemma holds for every word u and for every partial run V of $\mathcal{A}$ on u such that $height(V) < height(T)$. From $p \preceq_F r$, there is a transition $r \xrightarrow{w_1} R$ of $\mathcal{A}$ where such that $succ_T(p) \preceq_F^{\forall\exists} R$. Observe that $T = \{p\} \cup \bigcup_{p' \in succ_T(p)} pT(p')$, where for each $p' \in succ_T(p)$, $T(p')$ is a partial run of $\mathcal{A}$ with the root p' on the word v such that $w = w_1 v$. Notice that $height(T(p')) < height(T)$. The induction hypothesis now can be applied to every triple $p' \in succ_T(p), r' \in R, T(p')$ with $p' \preceq_F r'$. It gives us a partial run $U_{r'}$ of $\mathcal{A}$ on v with $root(U_{r'}) = r'$, such that $T(p') \preceq_F U_{r'}$. The run U with the required properties is then constructed by plugging the runs $U_{r'}, r' \in R$, to r, i.e., $U = \{r\} \cup \bigcup_{r' \in R} rU_{r'}$. ∎

PROOF. [Lemma 3] By induction on the length of π. In the base case, when $\pi = p$ and $T = \{p\}$, it is sufficient to take $U = \{r\}$ and $\psi = r$. Suppose now that $\pi \neq p$ and that the lemma holds for every partial run T' of $\mathcal{A}$ on w, states $p', r' \in Q$ such that $p' \preceq_B r'$, and every $\pi' \in branches(T')$ with $leaf(\pi') = p'$ and $|\pi'| < |\pi|$.

For the induction step, let $\pi = \pi' p$ and let $succ_T(\pi') = P \cup \{p\}, p \notin P$. By the definition of $\preceq_B$, there is a transition $s \xrightarrow{w_{|\pi|}} R \cup \{r\}, r \notin R$ of $\mathcal{A}$ such that $leaf(\pi') \preceq_B s$ and $P \preceq_F^{\forall\exists} R$. Let $T' = T \setminus \{\pi\} \setminus \bigcup_{p' \in P} \pi' T(\pi' p')$. Then T' is a partial run of $\mathcal{A}$ on w and $\pi' \in branches(T')$, $|\pi'| < |\pi|$, and therefore we can apply induction hypothesis to T', $leaf(\pi')$, s, and π'. This gives us a partial run U' of $\mathcal{A}$ on w with $\psi' \in branches(U')$ such that $leaf(\psi') = s$, $\pi' \preceq_B \psi'$ and for each $1 \leq j \leq |\pi'|$, $T' \ominus_j \pi' \preceq_F^{\forall\exists} U' \ominus_j \psi'$. For every $p' \in succ_T(\pi')$, $T(\pi' p')$ is a partial run of $\mathcal{A}$ with the root p' on the suffix v of w such that $w = uv, |u| = |\pi| - 1$. We can apply Lemma 2 to the triples $r' \in R, p' \in P, T(\pi' p')$ with $p' \preceq_F r'$. This gives us for each $r' \in R$ a run $U_{r'}$ of $\mathcal{A}$ on v with $root(U_{r'}) = r'$ such that there is some $p' \in P$ with $T(\pi' p') \preceq_F U_{r'}$. Now we construct a run U and a path ψ with the required properties by plugging r and runs $U_{r'}, r' \in R$ to the path ψ' in U', i.e., $\psi = \psi' r$ and $U = U' \cup \{\psi\} \cup \bigcup_{r' \in R} \psi' U_{r'}$. (To see that U really satisfies the required properties, observe the following: (i) As $U \ominus_{|\pi'|} \psi = \bigcup_{r' \in R} branches(U_{r'})$ and $T \ominus_{|\pi'|} \pi = \bigcup_{p' \in P} branches(T(\pi' p'))$, and because for each $r' \in R$, there is $p' \in P$ with $T(\pi' p') \preceq_F U_{r'}$, we have that $T \ominus_{|\pi'|} \pi \preceq_F^{\forall\exists} U \ominus_{|\pi'|} \psi$. (ii) For all $1 \leq j < |\pi'|$, $T \ominus_j \pi = T' \ominus_j \pi' \preceq_F^{\forall\exists} U' \ominus_j \psi' = U \ominus_j \psi$.). ∎

D Mediated Equivalence Can Be Used for Quotienting

We give full proofs of lemmas in Section 4.1 leading to Theorem 8 and Corollary 9.

D.1 Quotienting versus Extending

Lemma 12. $\mathcal{L}(\mathcal{A}/\equiv) \subseteq \mathcal{L}(\mathcal{A}^+)$.

Proof. Let $\mathcal{A}^+_\equiv = (\Sigma, Q, \iota, \delta^+_\equiv, \alpha^+_\equiv)$ be the automaton extended according to $\equiv$. Observe that states q and r with $q \equiv r$ are forward simulation equivalent in $\mathcal{A}^+_\equiv$. (q and r are in $\mathcal{A}^+_\equiv$ either both accepting or both nonaccepting, and for all $a \in \Sigma$, $\delta^+_\equiv(q,a) = \delta^+_\equiv(r,a)$). Gurumurthy et al. in [8] prove that quotienting with respect to forward simulation preserves language. Therefore, $\mathcal{L}(\mathcal{A}/\equiv) = \mathcal{L}(\mathcal{A}^+_\equiv)$. It is also easy to see that $\mathcal{L}(\mathcal{A}^+_\equiv) \subseteq \mathcal{L}(\mathcal{A}^+)$, as $\mathcal{A}^+$ has a richer transition function than $\mathcal{A}^+_\equiv$ and $\alpha^+ = \alpha^+_\equiv$. Thus, $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{A}^+)$. ∎

D.2 Relating Partial runs of $\mathcal{A}^+$ and $\mathcal{A}$

Proof. [Lemma 6] The proof of Lemma 6 relies on Lemma 3 and the definition of $\preceq_M$. We first choose a suitable branch π of $\mathsf{sw}_T(V)$ as follows. Let $1 \leq k \leq |\tau_V|$ be some index such that $\mathsf{sw}_T(V) \ominus_k \tau_V$ is nonempty. If $\tau_V = \epsilon$, then $k = 1$. We choose some $\pi' \in \mathsf{sw}_T(V) \ominus_k \tau_V$ which is minimal wrt. $\preceq_F$, meaning that there is no $\pi'' \in \mathsf{sw}_T(V) \ominus_k \tau_V$ different from π' such that $\pi'' \preceq_F \pi'$. We put $\pi = \tau_V^k \pi'$. We note that this is the place where we use the $\preceq_F$-unambiguity assumption. If $\mathcal{A}$ was $\preceq_F$-ambiguous, there need not be a k such that $\mathsf{sw}_T(V) \ominus_k \tau_V$ contains a minimal element wrt. $\preceq_F$.

From $ext(\pi) \preceq_M leaf(\pi)$, there is a mediator s with $ext(\pi) \preceq_F s \succeq_B leaf(\pi)$. We apply Lemma 3 to V, π, $leaf(\pi)$ and s, which give us a partial run W and $\psi \in branches(W)$ with $leaf(\psi) = s$ such that $\pi \preceq_B \psi$, and for all $1 \leq i \leq |\pi|$, $V \ominus_i \pi \preceq_F^{\forall\exists} W \ominus_i \psi$. Let $\tau_W = \psi$. The proof will be concluded by showing that (i) $T \preceq_{\mathsf{w}\text{-}ext} W$, (ii) $\tau_W \notin branches(\mathsf{sw}_T(W))$, and (iii) $\langle \tau_V, \mathsf{sw}_T(V)\rangle \sqsubset \langle \tau_W, \mathsf{sw}_T(W)\rangle$, which implies $V \prec^T_{\tau_V,\tau_W} W$.

(i) To show that $T \preceq_{\mathsf{w}\text{-}ext} W$, we proceed as follows. Observe that for every $\phi \in branches(W) \setminus \{\psi\}$ there is a branch $\phi' \in branches(V) \setminus \{\pi\}$ such that $leaf(\pi) \preceq_F leaf(\psi)$ and $\pi \preceq_\alpha \psi$. This holds because for all $1 \leq i \leq |\pi|$, $V \ominus_i \pi \preceq_F^{\forall\exists} W \ominus_i \psi$ and because $\pi \preceq_B \psi$ (to be more detailed, for every $\phi \in branches(W) \setminus \{\psi\}$, $\phi = \psi^i \rho$ for some i and $\rho \in W \ominus_i \psi$. There must be $\rho' \in V \ominus_i \pi$ with $\rho' \preceq_F \rho$. As $\pi \preceq_B \phi$, $\pi^i \preceq_B \phi^i$ which implies $\pi^i \preceq_\alpha \phi^i$. Similarly, $\rho' \preceq_F \rho$ implies $\rho' \preceq_\alpha \rho$ and also $leaf(\rho') \preceq_F leaf(\rho)$. Therefore, we can construct the branch $\pi^i\rho' \in branches(V) \setminus \{\pi\}$ with $\pi^i\rho' \preceq_\alpha \psi^i\rho = \phi$ and $leaf(\pi^i\rho') \preceq_F leaf(\psi^i\rho)$). We also know that $T \preceq_{\mathsf{w}\text{-}ext} V$, so $branches(T) \preceq^{\forall\exists}_{\mathsf{w}\text{-}ext} branches(V)$. Thus, by the definition of $\preceq_{\mathsf{w}\text{-}ext}$, we have that for every $\phi \in branches(W) \setminus \{\psi\}$, there are $\phi' \in branches(V)$ and $\phi'' \in branches(T)$ with $\phi'' \preceq_\alpha \phi' \preceq_\alpha \phi$ and $ext(\phi'') \preceq_M leaf(\phi') \preceq_F leaf(\phi)$. This by transitivity of α and the definition of $\preceq_M$ gives $\phi'' \preceq_\alpha \phi$ and $ext(\phi'') \preceq_M leaf(\phi)$, which means $\phi'' \preceq_{\mathsf{w}\text{-}ext} \phi$. As $T \preceq_{\mathsf{w}\text{-}ext} V$, there must also be a $\rho \in branches(T)$ with $\rho \preceq_{\mathsf{w}\text{-}ext} \pi$, and thus we have $\rho \preceq_\alpha \pi \preceq_B \psi$ and $ext(\rho) \preceq_F s = leaf(\psi)$, which by $\preceq_B \subseteq \preceq_\alpha$ and transitivity of $\preceq_\alpha$ gives $\rho \preceq_{ext} \psi$. As $\preceq_{ext} \subseteq \preceq_{\mathsf{w}\text{-}ext}$, this implies $\rho \preceq_{\mathsf{w}\text{-}ext} \psi$. Finally, from $root(T) \preceq_B root(V)$ (implied by $T \preceq_{\mathsf{w}\text{-}ext} V$), $\pi \preceq_B \psi$, and transitivity of $\preceq_B$, $root(T) \preceq_B root(W)$. This overall gives that $T \preceq_{\mathsf{w}\text{-}ext} W$.

(ii) Showing that $\psi \notin branches(\mathsf{sw}_T(W))$ is easy. In the above paragraph we have just shown that $\rho \preceq_{ext} \psi$ for some $\rho \in branches(T)$, this ψ is not a strict weakly covering branch.

(iii) To show that $\langle \tau_V, \mathsf{sw}_T(V)\rangle \sqsubset \langle \psi, \mathsf{sw}_T(W)\rangle$, we will proceed as follows: Will show that (a) for all $1 \leq i < k$, we have $\mathsf{sw}_T(V) \ominus_i \tau_V \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_i \psi$ and that (b) $\mathsf{sw}_T(V) \ominus_k \tau_V \prec_F^{\forall\exists}$

$\mathsf{sw}_T(W) \ominus_k \psi$. Notice first that for any partial run X of $\mathcal{A}$ and $\tau \in X$ with $\tau \notin \mathit{branches}(\mathsf{sw}_T(X))$, for all $1 \leq j \leq |\tau|$, $\mathsf{sw}_T(X) \ominus_j \tau \subseteq X \ominus_j \tau$. Recall that $\tau_V^k = \pi^k$, that $\mathsf{sw}_T(V) \ominus_k \tau_V$ is nonempty, and that for all $1 \leq i < |\pi|$, $V \ominus_i \pi \preceq_F^{\forall\exists} W \ominus_i \psi$.

We first show that for all $1 \leq i < |\pi|$, $\mathsf{sw}_T(V) \ominus_i \pi \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_i \psi$. For every $\phi \in \mathsf{sw}_T(W) \ominus_i \psi$, there is at least one $\phi' \in V \ominus_i \pi$ with $\phi' \preceq_F \phi$ (because $V \ominus_i \pi \preceq_F^{\forall\exists} W \ominus_i \psi$ and $\mathsf{sw}_T(W) \ominus_i \psi \subseteq W \ominus_i \psi$). We will show by contradiction that $\phi' \in \mathsf{sw}_T(V) \ominus_i \pi$ which will imply $\mathsf{sw}_T(V) \ominus_i \pi \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_i \psi$. Suppose that $\phi' \notin \mathsf{sw}_T(V) \ominus_i \pi$. Then the branch $\pi^i\phi'$ of V is not strict weakly covering, and as $T \preceq_{\text{w-}ext} V$, we have that there is some $\phi'' \in \mathit{branches}(T)$ with $\phi'' \preceq_{ext} \pi^i\phi'$. As $\pi \preceq_B \psi$, we have that $\pi^i \preceq_\alpha \psi^i$. As $\phi' \preceq_F \phi$, we have that $\phi' \preceq_\alpha \phi$ and $\mathit{leaf}(\phi') \preceq_F \mathit{leaf}(\phi)$. This together with $\phi'' \preceq_{ext} \pi^i\phi'$ gives that $\phi'' \preceq_\alpha \pi^i\phi' \preceq_\alpha \psi^i\phi$ and $\mathit{ext}(\phi'') \preceq_F \mathit{leaf}(\pi^i\phi') \preceq_F \mathit{leaf}(\psi^i\phi)$. By transitivity of $\preceq_\alpha$ and $\preceq_F$ and by the definition of $\preceq_{ext}$, we obtain $\phi'' \preceq_{ext} \psi^i\phi$. This contradicts with the fact that $\psi^i\phi$ is strict weakly covering (as $\phi \in \mathsf{sw}_T(W) \ominus_i \psi$) and therefore it must be the case that $\phi' \in \mathsf{sw}_T(V) \ominus_i \pi$.

(a) The fact that for all $1 \leq i < k$, $\mathsf{sw}_T(V) \ominus_i \tau_V \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_i \psi$ is implied by the result of the previous paragraph, because $\tau_V^k = \pi^k$ (thus $\mathsf{sw}_T(V) \ominus_i \tau_V = \mathsf{sw}_T(V) \ominus_i \pi$).

(b) It remains to show that $\mathsf{sw}_T(V) \ominus_k \tau_V \prec_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_k \psi$. By the definitions of $\ominus_k$, π and τ_V, it holds that $\mathsf{sw}_T(V) \ominus_k \pi \subset \mathsf{sw}_T(V) \ominus_k \tau_V$. (To see this, recall that π is strict weakly covering, but τ_V is not. Therefore, $\mathsf{sw}_T(V) \ominus_k \pi = \mathsf{sw}_T(V) \ominus_k \tau_V \setminus \mathit{branches}(\mathsf{sw}_T(V)(\pi^{k+1})))$. Thus, $\mathsf{sw}_T(V) \ominus_k \tau_V \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_k \psi$. As $\pi' \notin \mathsf{sw}_T(V) \ominus_k \pi$ and π' is a minimal element of $\mathsf{sw}_T(V) \ominus_k \tau_V$, $\mathsf{sw}_T(V) \ominus_k \pi \preceq_F^{\forall\exists} \mathsf{sw}_T(V) \ominus_k \tau_V$ cannot hold (there is no $\pi'' \in \mathsf{sw}_T(V) \ominus_k \pi$ with $\pi'' \preceq_F \pi'$). Therefore, $\mathsf{sw}_T(V) \ominus_k \tau_V \prec_F^{\forall\exists} \mathsf{sw}_T(V) \ominus_k \pi$, which together with $\mathsf{sw}_T(V) \ominus_k \pi \preceq_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_k \psi$ gives (by transitivity of $\preceq_F$) that $\mathsf{sw}_T(V) \ominus_k \tau_V \prec_F^{\forall\exists} \mathsf{sw}_T(W) \ominus_k \psi$. This completes the part (iii) of the proof and we can conclude that $V \prec_{\tau_V,\psi}^T W$.
∎

PROOF. [Lemma 7] The proof of Lemma 7 is done by induction to the structure of T, using Lemma 5 within the induction step. To make the induction argument pass, we will prove a stronger variant of the lemma. Let us first define the relation $\preceq_{\alpha^+\Rightarrow\alpha}^M$ on paths such that for two paths π and ψ, $\pi \preceq_{\alpha^+\Rightarrow\alpha}^M \psi$ iff $\pi \preceq_{\alpha^+\Rightarrow\alpha} \psi$ and $\mathit{leaf}(\pi) \preceq_M \mathit{leaf}(\psi)$. For two partial runs V and W, we use $V \preceq_{\alpha^+\Rightarrow\alpha}^M W$ to denote that $\mathit{branches}(V)$ $(\preceq_{\alpha^+\Rightarrow\alpha}^M)^{\forall\exists}$ $\mathit{branches}(W)$. Apparently, $\preceq_{\alpha^+\Rightarrow\alpha} \subseteq \preceq_{\alpha^+\Rightarrow\alpha}^M$ for paths as well ans for partial runs.

A stronger variant of the lemma: For any partial run T of $\mathcal{A}^+$ on $w \in \Sigma^\omega$, there exists a partial run U of $\mathcal{A}$ on w such that $\mathit{root}(T) \preceq_B \mathit{root}(U)$ and $T \preceq_{\alpha^+\Rightarrow\alpha}^M U$.

It is obvious that the above statement implies the statement of the lemma. We will prove it by induction to the structure of T. In the base case, $T = \{q\}$ for some $q \in Q$. If $q \notin \alpha^+$, we can put $U = \{q\}$ ($\preceq_M$ and $\preceq_B$ are reflexive). If $q \in \alpha^+$, then by the definition of α^+, there is $p \in \alpha$ such that $p \equiv_M q$. This means that $q \preceq_M p$ and $p \preceq_M q$. By the definition of $\preceq_M$, there exists a mediator s with $p \preceq_F s \succeq_B q$. As $\preceq_F \subseteq \preceq_\alpha$, $s \in \alpha$. Again by the definition of $\preceq_M$, $q \preceq_M p \preceq_F s \succeq_B q$ gives us $q \preceq_M s \succeq_B q$ and we can put $U = \{s\}$.

Suppose now that T is not only a root and that the stronger variant of the lemma holds for every partial run of $\mathcal{A}^+$ on w that is a proper subset of T. We choose some $\pi \in T$ such that $\mathit{succ}_T(\pi) \neq \emptyset$ and for every $p \in \mathit{succ}_T(\pi)$, $\mathit{succ}_T(\pi p) = \emptyset$. Denote $P = \mathit{succ}_T(\pi)$ and $q = \mathit{leaf}(\pi)$. Let $T' = T \setminus \{\pi p \mid p \in P\}$. T' is a partial run of $\mathcal{A}^+$ on w which is a proper subset of T, so we can apply the induction hypothesis. This gives us a partial run V of $\mathcal{A}$ on w such that $root(T') \preceq_B root(V)$ and $T' \preceq^M_{\alpha^+ \Rightarrow \alpha} V$.

Let $Bad_V \subseteq branches(V)$ be the set such that $\psi \in Bad_V$ iff there is no $\phi \in branches(T)$ such that $\phi \preceq^M_{\alpha^+ \Rightarrow \alpha} \psi$, and let $Good_V = branches(V) \setminus Bad_V$. Intuitively, Bad_V contains the problematic branches because of which $T \preceq^M_{\alpha^+ \Rightarrow \alpha} V$ does not hold.

By the definition of δ^+ and because $q \xrightarrow{w_{|\pi|}} P$ is a transition of $\mathcal{A}^+$, there must be some $s \in Q, s \preceq_M q$ where $s \xrightarrow{w_{|\pi|}} P$ is a transition of δ. We define an extension function ext_V such that $ext_V(\phi) = s$ for every $\phi \in Bad_V$ and $ext_V(\psi) = leaf(\psi)$ for every $\psi \in Good_V$. By applying Lemma 5 to V and ext_V, we get a partial run W of $\mathcal{A}$ on w with $V \preceq_{ext_V} W$. Now, for each $\psi \in branches(W)$, there is $\phi \in branches(V)$ with $\phi \preceq_{ext_V} \psi$. As $T' \preceq^M_{\alpha^+ \Rightarrow \alpha} V$, $\rho \preceq^M_{\alpha^+ \Rightarrow \alpha} \phi$ for some $\rho \in branches(T')$. There are two cases of how ρ and ψ may be related, depending on ϕ:

1. If $\phi \in Good_V$, then $ext(\phi) = leaf(\phi)$. In this case, by the definitions of $\preceq^M_{\alpha^+ \Rightarrow \alpha}$ and $\preceq_{ext_V}$, we have $\rho \preceq_{\alpha^+ \Rightarrow \alpha} \phi \preceq_\alpha \psi$ and $leaf(\rho) \preceq_M leaf(\phi) \preceq_F leaf(\psi)$, which gives $\rho \preceq_{\alpha^+ \Rightarrow \alpha} \psi$ and $leaf(\rho) \preceq_M leaf(\psi)$ (by the definition of $\preceq_M$), meaning that $\rho \preceq^M_{\alpha^+ \Rightarrow \alpha} \psi$.
2. To analyze the case when $\phi \in Bad_V$, observe that π is the only branch of T' which is not a branch of T. Therefore, it has to be the case that π is the only branch of T' with $\pi \preceq^M_{\alpha^+ \Rightarrow \alpha} \phi$ (If there was a another such a branch π' of T' with $\pi' \preceq^M_{\alpha^+ \Rightarrow \alpha} \phi$, then $\phi \in Good_V$ as $\pi' \in branches(T)$. There must be at leas one such a branch of T' as $T' \preceq^M_{\alpha^+ \Rightarrow \alpha} V$). Thus $\rho = \pi$. According to the definition of ext_V, $ext_V(\phi) = s$. Together with $V \preceq_{ext_V} W$, we have $\pi \preceq^M_{\alpha^+ \Rightarrow \alpha} \phi \preceq_\alpha \psi$ which gives $\pi \preceq_{\alpha^+ \Rightarrow \alpha} \psi$. However, we cannot guarantee any further relation between $leaf(\phi)$ and $leaf(\psi)$, and therefore we cannot derive $leaf(\pi) \preceq_M leaf(\psi)$ and $\pi \preceq^M_{\alpha^+ \Rightarrow \alpha} \psi$ as in the previous case.

We define the set $Bad_W \subseteq branches(W)$ such as $\psi \in Bad_W$ iff there is no $\rho \in T$ with $\rho \preceq^M_{\alpha^+ \Rightarrow \alpha} \psi$ and we let $Good_W = branches(W) \setminus Bad_V$. Analogically as Bad_V, Bad_W contains the branches because of which $T \preceq^M_{\alpha^+ \Rightarrow \alpha} W$ does not hold. Note that if $\psi \in Bad_V$, then all the $\phi \in branches(V)$ with $\phi \preceq_{ext_V} \psi$ are as in the case (2) above, i.e., π is the only branch of T' with $\pi \preceq^M_{\alpha^+ \Rightarrow \alpha} \phi$. By the definition of $\preceq_{ext_V}$, $s = ext_V(\phi) \preceq_F leaf(\psi)$. Therefore, by the definition of $\preceq_F$, there must be some transition $leaf(\psi) \xrightarrow{w_{|\pi|}} R_\psi$ of $\mathcal{A}$ where $P \preceq^{\forall\exists}_F R_\psi$. We extend W by firing these transitions for every $\psi \in Bad_W$, in which way we get a run $X = W \cup \{\psi R_\psi \mid \psi \in Bad_W\}$ of $\mathcal{A}$ on w.

Let us use $New_X = \{\psi R_\psi \mid \psi \in Bad_W\}$ to denote the branches of X that erased by firing the transitions. Observe that $branches(X) = Good_W \cup New_X$. Recall that for all $\psi \in Bad_W$, $\pi \preceq_{\alpha^+ \Rightarrow \alpha} \psi$ and that for every $\psi \in New_X$, there is some $p \in P$ such that $p \preceq_F leaf(\psi)$. We will define an extension function ext_X of X as follows:

1. If $\psi \in Good_W$, $ext_X(\psi) = leaf(\psi)$.
2. If $\psi \in New_X$ and there is $p \in P$ with $p \preceq_F leaf(\psi)$ and $p \preceq_{\alpha^+ \Rightarrow \alpha} leaf(\psi)$, we let $ext_X(\psi) = leaf(\psi)$.
3. If $\psi \in New_X$ and there is no $p \in P$ with $p \preceq_F leaf(\psi)$ and $p \preceq_{\alpha^+ \Rightarrow \alpha} leaf(\psi)$, we proceed as follwos. By the definition of New_X, there is some $p' \in P$ such that $p' \preceq_F leaf(\psi)$. Because of $\preceq_F \subseteq \preceq_\alpha$, the fact that $p' \preceq_F leaf(\psi)$ and not $p' \preceq_{\alpha^+ \Rightarrow \alpha} leaf(\psi)$ implies that $p', leaf(\psi) \notin \alpha$ and $p' \in \alpha^+$. This by the definition of α^+ means that there is some $v \in \alpha$ with $p' \equiv_M v$. We put $ext_X(\psi) = v$.

We apply Lemma 5 to X and ext_X, which gives us a partial run U of $\mathcal{A}$ on w with $X \preceq_{ext_X} U$. We will check that U satisfies the statement of the stronger variant of the lemma. We will first prove that that $T \preceq^M_{\alpha^+\Rightarrow\alpha} U$. For each $\tau \in branches(U)$, there is $\psi \in branches(X)$ with $\psi \preceq_{ext_X} \tau$. We will derive that there is some $\rho \in branches(T)$ with $\rho \preceq^M_{\alpha^+\Rightarrow\alpha} \tau$. The argument will depend on which of the above three types ψ is of:

1. If $\psi \in Good_W$, then there is some $\rho \in T$ with $\rho \preceq^M_{\alpha^+\Rightarrow\alpha} \psi$. Recall that $ext_X(\psi) = leaf(\psi)$ in this case. Thus, by the definitions of $\preceq^M_{\alpha^+\Rightarrow\alpha}$ and $\preceq_{ext_X}$, we have $\rho \preceq_{\alpha^+\Rightarrow\alpha} \psi \preceq_\alpha \tau$ and $leaf(\rho) \preceq_M leaf(\psi) \preceq_F leaf(\tau)$, which gives $\rho \preceq_{\alpha^+\Rightarrow\alpha} \tau$ and $leaf(\rho) \preceq_M leaf(\tau)$, i.e., $\rho \preceq^M_{\alpha^+\Rightarrow\alpha} \tau$.
2. If $\psi \in New_X$ and there is some $p \in P$ with $p \preceq_F leaf(\psi)$ and $p \preceq_{\alpha^+\Rightarrow\alpha} leaf(\psi)$, then by the definition of ext_X, $ext_X(\psi) = leaf(\psi)$. Recall that as $\psi^{|\psi|-1} \in Bad_W$, $\pi \preceq_{\alpha^+\Rightarrow\alpha} \psi^{|\psi|-1}$. Therefore, also $\pi p \preceq_{\alpha^+\Rightarrow\alpha} \psi$. By the definition of $\preceq_{ext_X}$, we have that $\psi \preceq_\alpha \tau$ and $leaf(\psi) \preceq_F leaf(\tau)$. Finally, $\pi p \preceq_{\alpha^+\Rightarrow\alpha} \psi \preceq_\alpha \tau$ and $p \preceq_F leaf(\psi) \preceq_F leaf(\tau)$ together imply that $\pi p \preceq^M_{\alpha^+\Rightarrow\alpha} \tau$.
3. If $\psi \in New_X$ and there is no $p \in P$ with $p \preceq_F leaf(\psi)$ and $p \preceq_{\alpha^+\Rightarrow\alpha} leaf(\psi)$, then by the definition of ext_X, $ext_X(\psi) = v$, where $v \in \alpha$ and $p' \equiv_M v, p' \preceq_F leaf(\psi)$ for some $p' \in P$. By $\psi \preceq_{ext_X} \tau$, we have $\psi \preceq_\alpha \tau$ and $v \preceq_F leaf(\tau)$. Thus, by the definition of $\preceq_M$, $p' \equiv_M v \preceq_F leaf(\tau)$ gives $p' \preceq_M leaf(\psi)$. As $\preceq_F \subseteq \preceq_\alpha$, we have that $leaf(\tau) \in \alpha$ and thus $p' \preceq_{\alpha^+\Rightarrow\alpha} leaf(\tau)$. As $\psi^{|\psi|-1} \in Bad_W$, we have that $\pi \preceq_{\alpha^+\Rightarrow\alpha} \psi^{|\psi|-1}$. Together with $\psi \preceq_\alpha \tau$, this gives $\pi p' \preceq_{\alpha^+\Rightarrow\alpha} \tau$. Therefore, $\pi p' \preceq^M_{\alpha^+\Rightarrow\alpha} \tau$.

Thus, we have proven that $T \preceq^M_{\alpha^+\Rightarrow\alpha} U$. Finally, as $V \preceq_{ext_V} W$ and $X \preceq_{ext_X} U$, we have $root(V) \preceq_B root(W)$ and $root(X) \preceq_B root(U)$. Together with $root(X) = root(W)$ and $root(T) = root(T') \preceq_B root(V)$, we have $root(T) \preceq_B root(V) \preceq_B root(X) \preceq_B root(U)$. By transitivity of $\preceq_B$, $root(T) \preceq_B root(U)$. We have verified that U satisfies the statement of the of the stronger variant of the lemma, which concludes the proof. ∎

D.3 Relating Accepting Runs of $\mathcal{A}^+$ and $\mathcal{A}$

LEMMA 13. *A run T of $\mathcal{A}$ with $root(T) = \iota$ is accepting if and only if for every $\pi \in T$, there exists a constant $k_\pi \in \mathbb{N}$ such that every ψ with $\pi\psi \in T$ and $|\psi| \geq k$ contains an accepting state.*

PROOF. *(if)* For every $\pi \in branches(T)$, there is an infinite sequence of $k_0, k_1 \ldots$ such that:

- $k_0 = 0$ and
- for all $i \in \mathbb{N}$, $k_i = k_{i-1} + k_{\pi^n}$ where $n = k_{i-1} + 1$.

For all $i \in \mathbb{N}$, every segment of π between $k_{i-1} + 1$ and k_i contains and accepting state, so, π contains infinitely many accepting states.

(only if) By contradiction. Suppose that there is $\pi \in T$ for which there is no k_π. We will show that in this case, there must be $\psi \in Q^\omega$ such that $\pi\psi \in branches(T)$ and ψ does not contain an accepting state (which contradicts with the assumption that T is accepting).

We will give a procedure which for each $i \in \mathbb{N}$ returns ψ^i (based on the knowledge of ψ^{i-1}). For each $i \in \mathbb{N}^0$, we will keep the invariant that for $\pi\psi^i$, there is no $k_{\pi\psi^i}$ and that ψ^i does not contain an accepting state. As $\psi^0 = \epsilon$, the invariant holds for $i = 0$.

Let the invariant hold for $i-1, i \in \mathbb{N}$, and suppose that we have already constructed ψ^{i-1}. There must be some $q \in succ_T(\pi\psi^{i-1})$ such that $q \notin \alpha$, otherwise $\pi\psi^{i-1} = 1$, which contradicts the invariant for $i-1$. We put $\psi^i = \psi^{i-1}q$. Observe that the invariant is satisfied for i too. Therefore, we can construct the ith prefix ψ^i of ψ that does not contain an accepting state for every $i \in \mathbb{N}$. This proves that there is a branch $\pi\psi$ of T where ψ does not contain an accepting state, which contradicts the assumption that T is accepting. ∎

LEMMA 14. *For every accepting run T of $\mathcal{A}^+$ a word $w \in \Sigma^\omega$, there exists an accepting run U of $\mathcal{A}$ on w.*

PROOF. For a tree X over Q, let $X(i) = \{\pi \in X \mid |\pi| \leq i\}$ be the ith prefix of X ($X(0) = \emptyset$). From Lemma 7, for each $i \in \mathbb{N}$, there is a partial run U_i of $\mathcal{A}$ on w such that $T(i) \preceq_{\alpha^+ \Rightarrow \alpha} U_i$ and $root(T(i)) \preceq_B root(U_i)$. As $\preceq_B \subseteq \preceq_\iota$, $root(U_i) = \iota$. Note that for all $\pi \in branches(U_i)$, $|\pi|$ equals i, because only paths of the same length can be related by $\preceq_{\alpha^+ \Rightarrow \alpha}$. Denote $\mathbb{U}^\infty = \{U_1, U_2, \ldots\}$. $\mathbb{U}^\infty$ is an infinite set that for each $k \in \mathbb{N}$ contains a partial run U_k of $\mathcal{A}$ with all the branches of the length k. We will use $\mathbb{U}^\infty$ to construct the infinite accepting run U.

Observe that for any infinite set $\mathbb{V}^\infty$ of partial runs of $\mathcal{A}$ and for any $i \in \mathbb{N}$, there has to be at least one partial run W of $\mathcal{A}$ such that for infinitely many $V \in \mathbb{V}^\infty$, $W = V(i)$. The reason is that for any $i \in \mathbb{N}$, there is obviously only finitely many of possible partial runs of the height i that $\mathcal{A}$ can generate.

We prove the existence of U by giving a procedure, which for every $k \in \mathbb{N}$ gives the kth prefix $U(k)$ of U.

- Let $\mathbb{U}_0^\infty = \mathbb{U}^\infty$ and let $U(0) = \emptyset$.
- For every $k \in \mathbb{N}$, $U(k)$ is derived from $U(k-1)$ as follows. Let $\mathbb{U}_k^\infty \subseteq \mathbb{U}^\infty$ be defined as the set such that for all $i \in \mathbb{N}$, $U_i \in \mathbb{U}_k^\infty$ iff $U(k-1) = U_i(k-1)$. In other words, $\mathbb{U}_k^\infty$ is the subset of $\mathbb{U}^\infty$ of the partial runs with the ith prefix equal to $U(k-1)$. Then, $U(k) = U_n(k)$ for some $n \geq k$ such that $U_n \in \mathbb{U}_k^\infty$ and there is infinitely many $m \in \mathbb{N}$ such that $U_m \in \mathbb{U}_k^\infty$ and $U_n(k) = U_m(k)$. I other words, $U(k)$ is a tree that appears as the kth prefix of infinitely many partial runs in $\mathbb{U}_k^\infty$.

To see that this construction is well defined, observe that:

- $\mathbb{U}_0^\infty$ is infinite, and
- for all $k \in \mathbb{N}$, if $\mathbb{U}_{k-1}^\infty$ is infinite, then $U(k-1)$ is defined and $\mathbb{U}_k^\infty$ is infinite.

Thus, $U(k)$ is well defined for every $k \in \mathbb{N}$ and U is a run of $\mathcal{A}$.

I remains to show that U is accepting. We will show that for every $\pi \in U$, there is $k_\pi \in \mathbb{N}$ such that every ψ with $\pi\psi \in T$ and $|\psi| \geq k$ contains an accepting state. By Lemma 13, it will follow that U is accepting.

Let us choose arbitrary $\pi \in U$. Let $n = |\pi|$. By Lemma 13, for every $\pi' \in branches(T_n)$, there is there is $k_{\pi'} \in \mathbb{N}$ such that every ψ' with $\pi'\psi' \in T$ and $|\psi'| \geq k_{\pi'}$ contains an accepting state. Let $k = \max\{k_{\pi'} \mid \pi' \in branches(T(n))\}$. By the construction of U, $T(n+k) \preceq_{\alpha^+ \Rightarrow \alpha} U(n+k)$. This implies that for every $\pi'' \in branches(U(n))$, every ψ'' with $\pi''\psi'' \in T$ and $|\psi''| \geq k$ contains an accepting state. As π in $branches(U(n))$, we can put $k_\pi = k$ and we are done. ∎

PROOF. [Theorem 8] The inclusion $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{A}^+)$ is obvious as $\mathcal{L}(\mathcal{A}^+)$ has riches both transition function and the set of accepting states. The inclusion $\mathcal{L}(\mathcal{A}^+) \subseteq \mathcal{L}(\mathcal{A})$ follows immediately from Lemma 14. ∎

E Correctness and Complexity of Computing $\preceq_B$

PROOF. [Theorem 10] *(if)* We define $\preceq$ to be a binary relation on Q such that $p \preceq r$ iff $p^{\odot} \preceq^{I} r^{\odot}$. We show that $\preceq$ is a backward simulation on Q which immediately implies the result.

Suppose that $p \preceq r$ and $p' \xrightarrow{a} \{p\} \cup P$ where $p \notin P$ is a transition of $\mathcal{A}$. Since $p \preceq r$, we know that $p^{\odot} \preceq^{I} r^{\odot}$; and since $p' \xrightarrow{a} \{p\} \cup P$ is a transition of $\mathcal{A}$, we know by definition of $A^{\odot}$ that $p^{\odot} \xrightarrow{a} (p', a, P)^{\odot}$ and $(p', a, P)^{\odot} \xrightarrow{a} p'^{\odot}$ are transitions in $A^{\odot}$. Since $\preceq^{I}$ is a simulation, we can find two transitions $r^{\odot} \xrightarrow{a} (r', a, R)^{\odot}$ and $(r', a, R)^{\odot} \xrightarrow{a} r'^{\odot}$ in $A^{\odot}$ with $(p', a, P)^{\odot} \preceq_L (r', a, R)^{\odot}$ and $p'^{\odot} \preceq_L r'^{\odot}$. From $p'^{\odot} \preceq^{I} r'^{\odot}$, $(p', a, P)^{\odot} \preceq^{I} (r', a, R)^{\odot}$, and the definition of the initial preorder I, we have $p' \preceq r'$ and $P \preceq_F^{\forall\exists} R$. It follows that $\preceq$ is in fact a backward simulation.

(only if) Define $\preceq_{\odot}$ as a binary relation on $Q^{\odot}$ such that $p^{\odot} \preceq_{\odot} r^{\odot}$ iff $p \preceq_B r$ and $(p, a, P)^{\odot} \preceq_{\odot} (r, a, R)^{\odot}$ iff $P \preceq_F^{\forall\exists} R$ and $p \preceq_B r$. By definition, $\preceq_{\odot} \subseteq I$. We show that $\preceq_{\odot}$ is a simulation on $Q^{\odot}$ which immediately implies the result. In the proof, we consider two sorts of states in $A^{\odot}$; namely those corresponding to states and those corresponding to "environments".

Suppose that $p^{\odot} \preceq_{\odot} r^{\odot}$ and the transition $p^{\odot} \xrightarrow{a} (p', a, P)^{\odot}$ is in $A^{\odot}$. Since $p^{\odot} \preceq_{\odot} r^{\odot}$, we know that $p \preceq_B r$. From the transition $p^{\odot} \xrightarrow{a} (p', a, P)^{\odot}$ and by definition of $A^{\odot}$, $p' \xrightarrow{a} P \cup \{p\}$ is a transition in $\mathcal{A}$. Since $p \preceq_B r$, there exists a transition $r' \xrightarrow{a} R \cup \{r\}$ in $\mathcal{A}$ such that $p' \preceq_B r'$ and $P \preceq_F^{\forall\exists} R$. It follows that there exists a transition $r^{\odot} \xrightarrow{a} (r', a, R)^{\odot}$ in $A^{\odot}$ such that $(p', a, P)^{\odot} \preceq_{\odot} (r', a, R)^{\odot}$.

Suppose that $(p, a, P)^{\odot} \preceq_{\odot} (r, a, R)^{\odot}$ and the transition $(p, a, P)^{\odot} \xrightarrow{a} p^{\odot}$ is in $A^{\odot}$. Since $(p, a, P)^{\odot} \preceq_{\odot} (r, a, R)^{\odot}$, we know that $P \preceq_F^{\forall\exists} R$ and $p \preceq_B r$. By definition of $A^{\odot}$, the transition $(r, a, R)^{\odot} \xrightarrow{a} r^{\odot}$ is in $A^{\odot}$. Since $p \preceq_B r$, we have $p^{\odot} \preceq_{\odot} r^{\odot}$. Together we have there exists a transition $(r, a, R)^{\odot} \xrightarrow{a} r^{\odot}$ in $A^{\odot}$ such that $p^{\odot} \preceq_{\odot} r^{\odot}$. It follows that $\preceq_{\odot}$ is a simulation on $Q^{\odot}$. ∎

PROOF. [Theorem 11] The complexity comes from three parts of the entire procedure: (1) compiling $\mathcal{A}$ into its corresponding LTS $A^{\odot}$, (2) computing the initial preorder I, and (3) running the algorithm for computing the LTS simulation. The LTS $A^{\odot}$ has at most nm+n states and $2nm$ transitions. It is trivial that Part (1) has both time and space complexity $O(nm)$. As we explained in Appendix A, Part(2) has time complexity $O(n^2m^2)$ and space complexity $O(nm^2)$. From the size of the LTS obtained from Part (1), Part (3) has time complexity $O(|\Sigma|n^2m^2)$ and space complexity $O(|\Sigma|n^2m^2)$. It follows that computing backward simulation has time complexity $O(|\Sigma|n^2m^2)$ and space complexity $O(|\Sigma|n^2m^2)$. Under our definition of ABA, every state has at least one outgoing transition for each symbol in Σ. It follows that $m \geq |\Sigma|n$. Therefore, we can also say that the procedure for computing maximal backward simulation has time complexity $O(nm^3)$ and space complexity $O(nm^3)$.

∎

F Counterexamples

F.1 Backward Simulation Cannot Be Used For Quotienting

Consider the following ABA $\mathcal{A} = (\{a,b\}, \{s_0,s_1,s_2,s_3,s_4,s_5,s_6\}, s_0, \delta, \{s_0,s_1,s_2,s_3,s_4,s_5,s_6\})$, where $s_0 \xrightarrow{a} \{s_4\}$, $s_0 \xrightarrow{a} \{s_1\}$, $s_0 \xrightarrow{b} \{s_0\}$, $s_1 \xrightarrow{b} \{s_2,s_5\}$, $s_1 \xrightarrow{b} \{s_1,s_3\}$, $s_2 \xrightarrow{b} \{s_2,s_3\}$, $s_3 \xrightarrow{a} \{s_0\}$, $s_4 \xrightarrow{b} \{s_4,s_6\}$, $s_5 \xrightarrow{b} \{s_0\}$, and $s_6 \xrightarrow{a} \{s_0\}$ are transitions of $\mathcal{A}$. The maximal forward simulation relation $\preceq_F$ in $\mathcal{A}$ is $\{(s_0,s_0), (s_1,s_0), (s_1,s_1), (s_1,s_5), (s_2,s_0), (s_2,s_1), (s_2,s_2), (s_2,s_4), (s_2,s_5), (s_3,s_3), (s_3,s_6), (s_4,s_0), (s_4,s_1), (s_4,s_2), (s_4,s_4), (s_4,s_5), (s_5,s_0), (s_5,s_5), (s_6,s_3), (s_6,s_6)\}$. The maximal backward simulation relation $\preceq_B$ parameterized with $\preceq_F$ is $\{(s_0,s_0), (s_1,s_1), (s_1,s_4), (s_2,s_2), (s_3,s_3), (s_4,s_1), (s_4,s_4), (s_5,s_2), (s_5,s_3), (s_5,s_5), (s_5,s_6), (s_6,s_2), (s_6,s_3), (s_6,s_5), (s_6,s_6)\}$.

If we collapse states wrt. $\preceq_M$ (i.e., two sets of states $\{s_1, s_4\}$, $\{s_5,s_6\}$ are collapsed), we will get the following ABA $\mathcal{A}' = (\{a,b\}, \{s_0,s_1,s_2,s_3,s_4\}, s_0, \delta, \{s_0,s_1,s_2,s_3,s_4\})$, where $s_0 \xrightarrow{a} \{s_1\}$, $s_0 \xrightarrow{b} \{s_0\}$, $s_1 \xrightarrow{b} \{s_2,s_4\}$, $s_1 \xrightarrow{b} \{s_1,s_4\}$, $s_1 \xrightarrow{b} \{s_1,s_3\}$, $s_2 \xrightarrow{b} \{s_2,s_3\}$, $s_3 \xrightarrow{a} \{s_0\}$, $s_4 \xrightarrow{a} \{s_0\}$, and $s_4 \xrightarrow{b} \{s_0\}$ are transitions of $\mathcal{A}'$.

Note that $\mathcal{A}'$ accepts the word ab^ω, but $\mathcal{A}$ does not.

F.2 Mediated Minimization Cannot Be Used On An $\preceq_F$-Ambiguous ABA

Consider the following ABA $\mathcal{A} = (\{a,b\}, \{s_0,s_1,s_2,s_3,s_4\}, s_0, \delta, \{s_4\})$, where $s_0 \xrightarrow{a} \{s_1,s_2,s_3\}$, $s_1 \xrightarrow{b} \{s_4\}$, $s_2 \xrightarrow{b} \{s_4\}$, $s_3 \xrightarrow{b} \{s_4\}$, $s_3 \xrightarrow{a} \{s_1,s_2,s_3\}$, and $s_4 \xrightarrow{a} \{s_4\}$ are transitions of $\mathcal{A}$. The maximal forward simulation relation $\preceq_F$ in $\mathcal{A}$ is $\{(s_0,s_0), (s_0,s_3), (s_1,s_1), (s_1,s_2), (s_1,s_3), (s_2,s_1), (s_2,s_2), (s_2,s_3), (s_3,s_3), (s_4,s_4)\}$. From $s_1 \equiv_F s_2$ and the transition $s_0 \xrightarrow{a} \{s_1,s_2,s_3\}$ we can find that $\mathcal{A}$ is $\preceq_F$-ambiguous. The maximal backward simulation relation $\preceq_B$ parameterized with $\preceq_F$ is $\{(s_0,s_0), (s_1,s_1), (s_1,s_2), (s_1,s_3), (s_2,s_1), (s_2,s_2), (s_2,s_3), (s_3,s_1), (s_3,s_2), (s_3,s_3), (s_4,s_4)\}$ and the mediated preorder $\preceq_M$ is $\{(s_0,s_0), (s_0,s_1), (s_0,s_2), (s_0,s_3), (s_1,s_1), (s_1,s_2), (s_1,s_3), (s_2,s_1), (s_2,s_2), (s_2,s_3), (s_3,s_1), (s_3,s_2), (s_3,s_3), (s_4,s_4)\}$.

If we collapse states wrt. $\preceq_M$ (i.e., merge the three states s_1, s_2, and s_3), we will get the following ABA $\mathcal{A}' = (\{a,b\}, \{s_0,s_1,s_2\}, s_0, \delta, s_2)$, where $s_0 \xrightarrow{a} \{s_1\}$, $s_1 \xrightarrow{a} \{s_1\}$, $s_1 \xrightarrow{b} \{s_2\}$, and $s_2 \xrightarrow{a} \{s_2\}$ are transitions of $\mathcal{A}'$. Note that $\mathcal{A}'$ accepts the word $aaba^\omega$, but $\mathcal{A}$ does not.